

EXPRESS MAIL MAILING LABEL NUMBER EL315197639US

PATENT  
10830.0078.NPUS00  
EMC-01-077

APPLICATION FOR UNITED STATES LETTERS PATENT

for

LARGEST MAGNITUDE INDICES SELECTION FOR (RUN, LEVEL)  
ENCODING OF A BLOCK CODED PICTURE

By

Ugur Sezer

Sorin Faibish

Seyfullah H. Oguz

Wayne W. Duso

## BACKGROUND OF THE INVENTION

#### 1. Field of the Invention

The present invention relates to processing of compressed visual data, and in particular the processing of compressed visual data in order to reduce data storage requirements or data transmission bandwidth at the expense of decreased quality.

## Background Art

7 It has become common practice to compress audio/visual data in order to reduce  
8 the capacity and bandwidth requirements for storage and transmission. One of the most  
9 popular audio/video compression techniques is MPEG. MPEG is an acronym for the  
10 Moving Picture Experts Group, which was set up by the International Standards  
11 Organization (ISO) to work on compression. MPEG provides a number of different  
12 variations (MPEG-1, MPEG-2, etc.) to suit different bandwidth and quality constraints.  
13 MPEG-2, for example, is especially suited to the storage and transmission of broadcast  
14 quality television programs.

15 For the video data, MPEG provides a high degree of compression (up to 200:1) by  
16 encoding 8 x 8 blocks of pixels into a set of discrete cosine transform (DCT) coefficients,  
17 quantizing and encoding the coefficients, and using motion compensation techniques to  
18 encode most video frames as predictions from or between other frames. In particular, the  
19 encoded MPEG video stream is comprised of a series of groups of pictures (GOPs), and  
20 each GOP begins with an independently encoded (intra) I frame and may include one or  
21 more following P frames and B frames. Each I frame can be decoded without  
22 information from any preceding and/or following frame. Decoding of a P frame requires  
23 information from a preceding frame in the GOP. Decoding of a B frame requires

1 information from both a preceding and a following frame in the GOP. To minimize  
2 decoder buffer requirements, transmission orders differ from presentation orders for some  
3 frames, so that all the information of the other frames required for decoding a B frame  
4 will arrive at the decoder before the B frame.

5 In addition to the motion compensation techniques for video compression, the  
6 MPEG standard provides a generic framework for combining one or more elementary  
7 streams of digital video and audio, as well as system data, into single or multiple program  
8 transport streams (TS) which are suitable for storage or transmission. The system data  
9 includes information about synchronization, random access, management of buffers to  
10 prevent overflow and underflow, and time stamps for video frames and audio packetized  
11 elementary stream packets embedded in video and audio elementary streams as well as  
12 program description, conditional access and network related information carried in other  
13 independent elementary streams. The standard specifies the organization of the  
14 elementary streams and the transport streams, and imposes constraints to enable  
15 synchronized decoding from the audio and video decoding buffers under various  
16 conditions.

17 The MPEG-2 standard is documented in ISO/IEC International Standard (IS)  
18 13818-1, "Information Technology-Generic Coding of Moving Pictures and Associated  
19 Audio Information: Systems," ISO/IEC IS 13818-2, "Information Technology-Generic  
20 Coding of Moving Pictures and Associated Audio Information: Video," and ISO/IEC IS  
21 13818-3, "Information Technology-Generic Coding of Moving Pictures and Associated  
22 Audio Information: Audio," which are incorporated herein by reference. A concise  
23 introduction to MPEG is given in "A guide to MPEG Fundamentals and Protocol

1 Analysis (Including DVB and ATSC)," Tektronix Inc., 1997, incorporated herein by  
2 reference.

3 MPEG-2 provides several optional techniques that allow video coding to be  
4 performed in such a way that the coded MPEG-2 stream can be decoded at more than one  
5 quality simultaneously. In this context, the word "quality" refers collectively to features  
6 of a video signal such as spatial resolution, frame rate, and signal-to-noise ratio (SNR)  
7 with respect to the original uncompressed video signal. These optional techniques are  
8 known as MPEG-2 scalability techniques. In the absence of the optional coding for such  
9 a scalability technique, the coded MPEG-2 stream is said to be nonscalable. The MPEG-  
10 2 scalability techniques are varieties of layered or hierarchical coding techniques, because  
11 the scalable coded MPEG-2 stream includes a base layer that can be decoded to provide  
12 low quality video, and one or more enhancement layers that can be decoded to provide  
13 additional information that can be used to enhance the quality of the video information  
14 decoded from the base layer. Such a layered coding approach is an improvement over a  
15 simulcast approach in which a coded bit stream for a low quality video is transmitted  
16 simultaneously with an independently coded bit stream for high quality video. The use of  
17 video information decoded from the base layer for reconstructing the high quality video  
18 permits the scalable coded MPEG-2 stream to have a reduced bit rate and data storage  
19 requirement than a comparable simulcast data stream.

20 The MPEG-2 scalability techniques are useful for addressing a variety of  
21 applications, some of which do not need the high quality video that can be decoded from  
22 a nonscalable coded MPEG stream. For example, applications such as video  
23 conferencing, video database browsing, and windowed video on computer workstations

1 do not need the high quality provided by a nonscalable coded MPEG-2 stream. For  
2 applications where the high quality video is not needed, the ability to receive, store, and  
3 decode an MPEG-2 base-layer stream having a reduced bit rate or data storage capacity  
4 may provide a more efficient bandwidth versus quality tradeoff, and a more efficient  
5 complexity versus quality tradeoff. A scalable coded MPEG-2 stream provides  
6 compatibility for a variety of decoders and services. For example, a reduced complexity  
7 decoder for standard television could decode a scalable coded MPEG-2 stream produced  
8 for high definition television. Moreover, the base layer can be coded for enhanced error  
9 resilience and can provide video at reduced-quality when the error rate is high enough to  
10 preclude decoding at high quality.

11 The MPEG scaling techniques are set out in sections 7.7 to 7.11 of the MPEG-2  
12 standard video encoding chapter 13818-2. They are further explained in Barry G.  
13 Haskell et al., Digital Video: An Introduction to MPEG-2, Chapter 9, entitled “MPEG-2  
14 Scalability Techniques,” pp. 183-229, Chapman & Hall, International Thomson  
15 Publishing, New York, 1997, incorporated herein by reference. The MPEG scalability  
16 techniques include four basic techniques, and a hybrid technique that combines at least  
17 two of the four basic techniques. The four basic techniques are called data partitioning,  
18 signal-to-noise ratio (SNR) scalability, spatial scalability, and temporal scalability.

19 Data partitioning is a method of partitioning a single layer coded bit-stream into  
20 two classes, including a base layer “partition 0” and an enhancement layer “partition 1”.  
21 Partition 0 contains all high level header information as well as some low frequency  
22 discrete cosine transform (DCT) coefficients. Partition 1 contains all remaining higher  
23 frequency DCT coefficients and end-of-block (EOB) markers. Some syntax elements

1 belonging to partition 0 are redundantly copied to partition 1 to facilitate error recovery.  
2 This duplicated information includes the sequence\_header, GOP\_header, picture\_header,  
3 sequence\_end\_code, sequence\_extension, picture\_extension, and  
4 sequence\_scalable\_extension. This duplication ensures that there is proper  
5 synchronization and recovery following a bit-stream error in the low priority  
6 enhancement layer (partition 1) and introduces very little overhead. With respect to the  
7 single layer coded bit-stream, the separation point between the syntax elements to be  
8 included in the base and enhancement layers is indicated by a priority breakpoint (PBP)  
9 marker. The PBP can be adjusted at every picture slice. The PBP marker partitioning  
10 granularity is at the (run, level) DCT event level of the coded block data. Data  
11 partitioning is especially useful for error resilient video transmission over asynchronous  
12 transfer mode (ATM) networks and other networks where data prioritization is possible.  
13 Data partitioning has a number of shortcomings, including limited flexibility for PBP  
14 adjustment (in terms of partitioning granularity and update frequency), and the  
15 accumulation of drift errors over P pictures due to partially available coefficient  
16 information from a damaged enhancement layer.

17 SNR scalability is a method of generating a multiplex of bit-streams representing  
18 individual layers including a base layer which contains DCT coefficients quantized at a  
19 basic moderate quality level, and one or more SNR enhancement layers that contain DCT  
20 refinement coefficients intended to enhance the precision of quantized DCT coefficients  
21 reconstructed based on the content of all lower layers. Consequently, SNR scalability is  
22 also referred to as “Quantization Noise Scalability.” The layers in SNR scalability are all  
23 at the same spatial and temporal resolutions but cumulatively produce increasing quality

1 levels starting with the lowest quality at the base layer. The base layer includes all high  
2 level header information, all motion compensation and macroblock (MB) type  
3 information, and coarse quantized DCT coefficient information. The enhancement layers  
4 include quantized DCT refinement coefficient information, and some amount of overhead  
5 information. The slice structure should be the same for all layers. Use of different  
6 quantization matrices in the base and enhancement layers is allowed. The overhead  
7 required by SNR scalability results in a decreased bandwidth utilization efficiency  
8 compared to data partitioning. SNR scalability is especially useful for simultaneous  
9 distribution of standard definition television and high-definition television, error-resilient  
10 video services over ATM and other networks, and multi-quality Video On Demand  
11 (VOD) services. SNR scalability has a number of shortcomings, including increased  
12 complexity and overhead as compared to data partitioning, inflexibility in bandwidth  
13 distribution among the layers primarily due to the fact that all motion information has to  
14 be carried in the base layer, and the shortcoming that no single SNR scalable codec can  
15 eliminate drift errors and also be reliable under lossy enhancement layer transmission.

16 There are two variations to SNR scalability, namely, chroma simulcast and  
17 frequency domain SNR (FDSNR) scalability. Chroma simulcast provides a means for  
18 simultaneous distribution of video services that use 4:2:0 and 4:2:2 chroma subsampling  
19 formats. The associated bit-stream structure has three layers, including a base layer, an  
20 enhancement layer, and a simulcast layer. The base layer is a distribution of video in the  
21 4:2:0 format. The enhancement layer provides SNR enhancement for the luminance  
22 component of the base layer. The simulcast layer includes chrominance components of  
23 the 4:2:2 format.

1 Frequency domain SNR scalability provides a transform domain method to  
2 achieve spatial resolution scalability. The base layer is intended for display at reduced  
3 spatial resolution and includes video encoded by a quantization matrix that allows a  
4 proper subset of normal size DCT transform coefficients to be selected and included in  
5 the base layer for use in conjunction with a smaller size DCT at the base layer decoder.  
6 The enhancement layer is the set of remaining normal size DCT transform coefficients.

7 Spatial scalability provides an ability to decode video at different spatial  
8 resolutions without first having to decode an entire (full-size) frame and then decimating  
9 it. The base layer carries the lowest spatial resolution version of the video obtained by  
10 decimating the original (full-size) video. Enhancement layers carry the differential  
11 information required to generate successively higher spatial resolution versions of the  
12 video. Spatial scalability supports interoperability between different video resolution and  
13 formats, such as support for simultaneous transmission of high definition television and  
14 standard definition television, and backward compatibility of MPEG-2 with different  
15 standards such as H.262 or MPEG-1. Spatial scalability supports error-resilient video  
16 transmission on ATM and other networks. Decoder complexity can scale with channel  
17 bandwidth. Spatial scalability has the advantages of a high degree of flexibility in video  
18 resolution and formats to be used for each layer, and a high degree of flexibility in  
19 achieving bandwidth partitioning between layers. There are no decoder drift problems  
20 because there are independent coding loops that are only loosely coupled. Spatial  
21 scalability, however, requires significantly increased complexity as compared to data  
22 partitioning and SNR scalability.

1 Temporal scalability provides an ability to decode video at different frame rates  
2 without first having to decode every single frame. The base layer carries the lowest  
3 frame rate version of the video coded by itself at the basic temporal rate. This version of  
4 the video is obtained from the original full frame rate version by a temporal down-  
5 sampling operation. The enhancement layers carry the information to construct the  
6 additional frames required to generate successively higher temporal resolution versions of  
7 the video. Additional frames in each enhancement layer are coded with temporal  
8 prediction relative to the frames carried by lower layers. Temporal scalability provides  
9 simultaneous support for different frame rates in the form of downward compatibility  
10 with lower-rate services, such as migration from first generation interlaced high  
11 definition television to high temporal resolution progressive high-definition television.  
12 Temporal scalability supports error-resilient video transmission on ATM and other  
13 networks. Decoder complexity can scale with channel bandwidth. Temporal scalability  
14 has the advantages of providing flexibility in achieving bandwidth partitioning between  
15 layers. There are no decoder drift problems because there are independent coding loops  
16 that are only loosely coupled. Temporal scalability has less complexity and higher  
17 efficiency than spatial scalability. Temporal scalability, however, provides a bandwidth  
18 partitioning flexibility that is more limited than spatial scalability because temporal  
19 scalability uses the same spatial resolution in all layers.

20 Hybrid scalability combines two scalabilities at a time from among SNR, spatial  
21 and temporal scalabilities. A base layer carries a basic quality, spatial and temporal  
22 resolution version of the intended video content. A first enhancement layer carries  
23 differential information required to implement one of the two intended enhancements on

1 the base layer. A second enhancement layer carries differential information required to  
2 implement the second intended enhancement on the combination of the base and the first  
3 enhancement layers. Hybrid scalability is useful in more demanding applications  
4 requiring scalability in two video quality aspects within three or more bit-stream layers.

5

6 **SUMMARY OF THE INVENTION**

7 In accordance with one aspect of the present invention, there is provided a method  
8 of producing a (run, level) encoded picture from an original picture. The method  
9 includes producing respective sets of transform coefficients for blocks of pixels in the  
10 original picture. Transform coefficients in the respective sets of transform coefficients  
11 are quantized to produce respective sets of quantization indices for the blocks of pixels.  
12 The quantization indices for at least some of the blocks are produced by using a  
13 quantization step size that is not uniform among at least some of the blocks. The method  
14 further includes selecting largest magnitude quantization indices from the respective sets  
15 of quantization indices to produce respective sets of quantization indices for the blocks of  
16 pixels. The method further includes (run, level) encoding quantization indices from the  
17 respective sets of quantization indices to produce the (run, level) encoded picture.

18 In accordance with another aspect, the invention provides a method of scaling  
19 non-scalable MPEG-2 coded video to produce reduced-bandwidth, reduced-quality  
20 MPEG-2 coded video. The non-scalable MPEG-2 coded video includes a set of non-zero  
21 AC discrete cosine transform (DCT) coefficients for 8x8 blocks of the non-scalable  
22 MPEG-2 coded video. The method includes removing non-zero AC DCT coefficients  
23 from the non-scalable MPEG-2 coded video so that the reduced-quality MPEG-2 coded

1 video includes no more than a selected number of largest magnitude quantization indices  
2 for the non-zero AC DCT coefficients for each 8x8 block.

3 In accordance with yet another aspect, the invention provides a digital computer  
4 for producing a (run, level) encoded picture from an original picture. The digital  
5 computer includes at least one processor programmed for producing respective sets of  
6 transform coefficients for blocks of pixels in the original picture, and quantizing  
7 transform coefficients in the respective sets of transform coefficients to produce  
8 respective sets of quantization indices for the blocks of pixels. Quantization indices for  
9 at least some of the blocks are produced by using a quantization step size that is not  
10 uniform among at least some of the blocks. The processor is further programmed for  
11 selecting largest magnitude quantization indices from the respective sets of quantization  
12 indices to produce respective sets of quantization indices for the blocks of pixels; and  
13 (run, level) encoding quantization indices from the respective sets of quantization indices  
14 to produce the (run, level) encoded picture.

15 In accordance with still another aspect, the invention provides a digital computer  
16 for scaling non-scalable MPEG-2 coded video to produce reduced-bandwidth, reduced-  
17 quality MPEG-2 coded video. The non-scalable MPEG-2 coded video including a set of  
18 non-zero AC discrete cosine transform (DCT) coefficients for 8x8 blocks of the non-  
19 scalable MPEG-2 coded video, the digital computer including a processor programmed  
20 for removing non-zero AC DCT coefficients from the non-scalable MPEG-2 coded video  
21 so that the reduced-quality MPEG-2 coded video includes no more than a selected  
22 number of largest magnitude quantization indices for the non-zero AC DCT coefficients  
23 for each 8x8 block.

## BRIEF DESCRIPTION OF THE DRAWINGS

2           Other objects and advantages of the invention will become apparent upon reading  
3       the following detailed description with reference to the accompanying drawings, in  
4       which:

5 FIG. 1 is a block diagram of a data network including a video file server  
6 implementing various aspects of the present invention;

7 FIG. 2 is a flowchart of a procedure executed by a stream server computer in the  
8 video file server of FIG. 1 to service client requests;

9 FIG. 3 is a flowchart of a procedure for splicing MPEG clips;

FIG. 4 is a flowchart of a procedure for seamless video splicing of MPEG clips;

11 FIG. 5 is a more detailed flowchart of the procedure for seamless video splicing  
12 of MPEG clips;

13 FIG. 6 is a continuation of the flowchart begun in FIG. 5;

14 FIG. 7 is a timing diagram showing a timing relationship between video  
15 presentation units (VPUs) and associated audio presentation units (APUs) in an original  
16 MPEG-2 coded data stream;

17 FIG. 8 is a timing diagram showing a timing relationship between video  
18 presentation units (VPUs) and associated audio presentation units (APUs) for a fast-  
19 forward trick-mode stream;

20 FIG. 9 is a flowchart of a procedure for selection and alignment of audio  
21 presentation units (APUs) in the fast-forward trick-mode stream;

22 FIG. 10 is a flowchart of a procedure for producing a trick-mode MPEG-2  
23 transport stream from a regular MPEG-2 transport stream (TS);

1 FIG. 11 is a diagram illustrating relationships between the MPEG discrete cosine  
2 transform (DCT) coefficients, spatial frequency, and the typical zig-zag scan order;

3 FIG. 12 is a diagram illustrating a relationship between an MPEG-2 coded bit  
4 stream and a reduced-quality MPEG-2 coded bit stream resulting from truncation of high-  
5 order DCT coefficients;

6 FIG. 13 is a flowchart of a procedure for scaling MPEG-2 coded video using a  
7 variety of techniques;

8 FIG. 14 is a flowchart of a procedure for signal-to-noise ratio scaling MPEG-2  
9 coded video using a frequency-domain low-pass truncation (FDSNR\_LP) technique;

10 FIG. 15 is a flowchart of a procedure for signal-to-noise ratio scaling MPEG-2  
11 coded video using a frequency-domain largest-magnitude coefficient selection  
12 (FDSNR\_LM) technique;

13 FIG. 16 is a flowchart of a procedure that selects one of a number of techniques  
14 for finding a certain number "k" of largest values out of a set of "n" values;

15 FIG. 17 is a flowchart of a procedure for finding a certain number "k" of largest  
16 values from a set of "n" values, which is used in the procedure of FIG. 16 for the case of  
17  $k < \frac{1}{2} n$ ;

18 FIG. 18 is a diagram of a hash table and associated hash lists;

19 FIG. 19 is a flowchart of a procedure for finding a certain number "k" of values  
20 that are not less than the smallest of the "k" largest values in a set of "n" values beyond a  
21 certain amount.

1 FIG. 20 is a flowchart of modification of the procedure of FIG. 15 in order to  
2 possibly eliminate escape sequences in the (run, level) coding of the largest magnitude  
3 coefficients;

4 FIG. 21 is a flowchart of a subroutine called in the flowchart of FIG. 20 in order  
5 to possibly eliminate an escape sequence;

6 FIG. 22 is a first portion of a flowchart of a procedure for scaling an MPEG-2  
7 coded video data stream using the modified procedure of FIG. 20 while adjusting the  
8 parameter "k" to achieve a desired bit rate, and adjusting a quantization scaling factor  
9 (QSF) to achieve a desired frequency of occurrence of escape sequences;

10 FIG. 23 is a second portion of the flowchart begun in FIG. 22;

11 FIG. 24 is a simplified block diagram of a volume containing a main file, a  
12 corresponding fast forward file for trick mode operation, and a corresponding fast reverse  
13 file for trick mode operation;

14 FIG. 25 is a more detailed block diagram of the volume introduced in FIG. 24;

15 FIG. 26A is a diagram showing video file access during a sequence of video  
16 operations including transitions between the main file, the related fast forward file, and  
17 the related fast reverse file;

18 FIG. 26B shows a script of a video command sequence producing the sequence of  
19 video play shown in FIG. 26A;

20 FIG. 27 is a table of read and write access operations upon the volume of FIG. 24  
21 and access modes that are used for the read and write access operations;

22 FIG. 28 is a hierarchy of video service classes associated with the fast forward file  
23 and the fast reverse file in the volume of FIG. 25;

1 FIG. 29 shows a system for modifying and combining an MPEG-2 audio-visual  
2 transport stream with an MPEG-2 closed-captioning transport stream to produce a  
3 multiplexed MPEG-2 transport stream having the same bit rate as the original MPEG-2  
4 audio-visual transport stream;

5 FIG. 30 shows a flowchart of a procedure for signal-to-noise ratio scaling MPEG-  
6 2 coded video using a frequency-domain largest magnitude indices selection  
7 (FDSNR\_LMIS) technique;

8 FIG. 31 shows a graph of the picture signal-to-noise ratio (PSNR) as a function of  
9 the number of bits used for only AC coefficients' encoding using the largest magnitude  
10 coefficient selection (LMCS) and largest magnitude indices selection (LMIS) procedures  
11 for quantization scale values (qsv) of 2, 4, 6, and 8, without the insertion of pivots;

12 FIG. 32 shows a graph of the picture signal-to-noise ratio (PSNR) as a function of  
13 the number of bits used for only AC coefficients' encoding using the largest magnitude  
14 coefficient selection (LMCS) and largest magnitude indices selection (LMIS) procedures  
15 for quantization scale values (qsv) of 12, 16, 20, and 24, without the insertion of pivots;

16 FIG. 33 shows a flowchart showing the successive application of a Pivot-1  
17 technique, a Pivot-2 technique, and a Pivot 3 technique for selection or insertion of pivot  
18 indices in order to avoid escape sequences or reduce the number of bits for (run, level)  
19 encoding;

20 FIG. 34 shows a graph of the average number of escape sequences per frame and  
21 a function of the number of AC coefficients retained in each block for a quantization  
22 scale value (qsv) of four for largest magnitude coefficient selection (LMCS) for no pivot  
23 insertion and for pivot insertion by each of the Pivot-1, Pivot-2, and Pivot-3 techniques;

1 FIG. 35 shows a graph of the average number of escape sequences per frame and  
2 a function of the number of AC coefficients retained in each block for a quantization  
3 scale value (qsv) of twenty-four for largest magnitude coefficient selection (LMCS) for  
4 no pivot insertion and for pivot insertion by each of the Pivot-1, Pivot-2, and Pivot-3  
5 techniques;

6 FIG. 36 shows a series of coefficients in a scan order, in order to illustrate the run  
7 length for a non-zero AC coefficient, and the insertion of a pivot index to reduce the run  
8 length;

9 FIG. 37 shows a pivot table indicating whether or not a pivot should be inserted  
10 for a given run length and level magnitude;

11 FIG. 38 shows a first sheet of a flowchart of a specific implementation of pivot  
12 insertion for avoiding escape sequences;

13 FIG. 39 is a second sheet of the flowchart begun in FIG. 38;

14 FIG. 40 is a flowchart of a procedure for a lookup in the pivot table of FIG. 37;

15 FIG. 41 is a flow diagram showing an encoding and decoding sequence including  
16 the insertion of noise in the form of pivot indices during encoding to reduce the number  
17 of bits in (run, level) encoding, and partial removal of the noise during decoding;

18 FIG. 42 is a flowchart showing how the process of pivot insertion during  
19 encoding or transcoding may be different depending on whether or not the decoder will  
20 attempt removal of the pivots;

21 FIG. 43 is a flowchart showing the removal of pivots during decoding;

1 FIG. 44 is a flowchart showing how the decoder determines whether or not a  
2 coefficient is possibly a pivot and whether or not a coefficient that is possibly a pivot is  
3 likely to be a pivot; and

4 FIG. 45 is a flow diagram showing the use of pivot insertion with transcoding or  
5 encoding with the process of largest magnitude indices selection (LMIS) or largest  
6 magnitude coefficient selection (LMCS).

7 While the invention is susceptible to various modifications and alternative forms,  
8 specific embodiments thereof have been shown by way of example in the drawings and  
9 will be described in detail. It should be understood, however, that it is not intended to  
10 limit the form of the invention to the particular forms shown, but on the contrary, the  
11 intention is to cover all modifications, equivalents, and alternatives falling within the  
12 scope of the invention as defined by the appended claims.

13

14                   DESCRIPTION OF ILLUSTRATIVE EMBODIMENTS

15                   I. Applications for Efficient Scaling of Non-Scalable MPEG-2 Video

16         With reference to FIG. 1, there is shown a block diagram of a data network 20  
17 linking a number of clients 21, 22, 23 to a video file server 24 implementing various  
18 aspects of the present invention. The video file server 24 includes at least one stream  
19 server computer 25 and a data storage system 26. The stream server computer 25 has a  
20 processor 27 and a network link adapter 28 interfacing the processor to the data network  
21 20. The processor 27 executes a data streaming program 29 in memory 30 in order to  
22 stream MPEG coded video in real-time to the clients.

1 Client requests for real-time video are placed in client play lists 31 in order to  
2 schedule in advance video file server resources for the real-time streaming of the MPEG  
3 coded video. The play lists 31 specify a sequence of video clips, which are segments of  
4 MPEG-2 files 32, 33 in data storage 34 of the data storage system 26. The stream server  
5 processor 27 accesses a client play list in advance of the time to begin streaming MPEG  
6 coded video from a clip, and sends a video prefetch command to a storage controller 35  
7 in the data storage system 26. The storage controller responds to the video prefetch  
8 command by accessing the clip in the data storage 34 to transfer a segment of the clip to  
9 cache memory 36. When the video data of the segment needs to be sent to the client, the  
10 stream server processor 27 requests the data from the storage controller 35, and the  
11 storage controller immediately provides the video data from the cache memory 36.  
12 Further details regarding a preferred construction and programming of the video file  
13 server 24 are disclosed in Duso et al., U.S. Patent 5,892,915 issued Apr. 6, 1999, entitled  
14 "System Having Client Sending Edit Commands to Server During Transmission Of  
15 Continuous Media From One Clip in Play List for Editing the Play List," incorporated  
16 herein by reference.

17 In accordance with an aspect of the invention, the stream server computer 25  
18 executes an MPEG scaling program 38 to produce reduced-quality MPEG coded video  
19 from nonscalable MPEG-2 coded video by truncating discrete cosine transform (DCT)  
20 AC coefficients from the coded blocks in the MPEG-2 coded video data. The reduced-  
21 quality MPEG coded video can be produced during ingestion of an MPEG-2 file 32 from  
22 the network 20, and stored in one or more associated files 37. Alternatively, the reduced-  
23 quality MPEG coded video in the files 37 could be produced as a background task from

1 the MPEG-2 file 32. Reduced-quality MPEG coded video could also be produced in  
2 real-time from an MPEG-2 file 33 during streaming of the reduced-quality MPEG coded  
3 video from the stream server computer 25 to the network 20. The reduced-quality MPEG  
4 coded video is useful for a variety of applications, such as browsing and review of stored  
5 MPEG-2 assets for search and play-list generation, bit stream scaling for splicing, and  
6 bit-rate adjustment via video quality alteration for services with limited resources.

7 A typical example of browsing for play-list generation involves searching stored  
8 assets in a multi-media data base for segments of a desired content to be included in the  
9 play list, and in particular selecting the beginning frame and ending frame of each  
10 segment to be included. Such editing occurs often in the broadcast environment for  
11 inserting commercials and news clips into pre-recorded television programming, and for  
12 editing movies for content and time compression. The decoding technique of the present  
13 invention permits a PC workstation 23 to perform the decoding and display in real-time  
14 by execution of a software program. An operator can view the video content in a display  
15 window 39 in a fast-forward or fast-reverse mode, stop at and resume from freeze frames  
16 that are valid “in points” and “out points” for seamless splicing, and select an in-point  
17 and out-point for a next segment to be included in the play list. The stream server  
18 computer 25 could also include a seamless splicing program 40 providing seamless  
19 transitions between video segments that are contiguous in a play list and are from  
20 different video clips.

21 For seamless splicing, it is often necessary to reduce the bitrate for one or more  
22 frames at the end of a first segment prior to splicing to a second segment. In this case the  
23 bitrate must be reduced to avoid buffer overflow as a result of displaying the original

1 frames at the end of the first segment. One method of reducing the bitrate is to insert a  
2 freeze frame at the end of the first segment, but this has the disadvantage of introducing  
3 distortion in the temporal presentation of the frames and precluding frame accuracy. A  
4 less disruptive method is to use the present invention for reducing the bitrate for a lower-  
5 quality presentation of one or more frames at the end of the first segment.

6 The present invention can also reduce the transmission bit rate and storage  
7 requirements for MPEG-2 applications by altering the video quality. For example,  
8 different clients may present different bandwidth access requests for video from  
9 nonscalable MPEG-2 files 32, 33 in the video file server. Also, temporary network  
10 congestion may limit the bandwidth available to satisfy a request for real-time streaming  
11 of video data. In each case, the present invention can alter the video quality to meet the  
12 desired or available bandwidth to satisfy the request.

13 With reference to FIG. 2, there is shown a flowchart of a procedure executed by a  
14 stream server computer in the video file server of FIG. 1 to service client requests. In a  
15 first step 50, execution branches to step 51 when a client request is not a request for real-  
16 time streaming. If the request is a request to input a new MPEG-2 file, then execution  
17 branches to step 52 to input the new MPEG-2 file and to create a reduced-quality version  
18 of the MPEG-2 file as available resources permit. If the request is not a request to input a  
19 new MPEG-2 file, then execution continues from step 51 to step 53. In step 53,  
20 execution branches to step 54 if the request is for play list editing. In step 54, the client  
21 may browse through the reduced-quality MPEG file to select in-points and out-points of  
22 clips to be spliced.

1        In step 50, when the request is for real-time streaming, then execution branches to  
2        step 55. In step 55, if there is network congestion so that there is insufficient bandwidth  
3        to transmit a stream of original-quality MPEG-2 coded video, then execution branches to  
4        step 56 to stream compressed video from the reduced-quality MPEG file. If no reduced-  
5        quality MPEG file is available for the desired clip, then the reduced-quality MPEG coded  
6        video to be streamed is produced in real-time from the original-quality MPEG-2 coded  
7        video. There are also applications, such as the display of spatially down-sampled video  
8        in a small display window (39 in FIG. 1), for which the client may request reduced-  
9        quality MPEG coded video. In this case, in the absence of network congestion, execution  
10      will continue from step 55 to step 57, and branch from step 57 to step 56 for streaming of  
11      reduced-quality MPEG coded video to the client.

12       Reduced-quality MPEG coded video is also useful for “trick-mode” operation.  
13       Trick-mode refers to fast forward or fast reverse display of video, in a fashion analogous  
14      to the fast forward and fast reverse playback functions of a video cassette recorder  
15      (VCR). The problem with trick-mode operation is that the transmission rate of the  
16      MPEG stream cannot simply be increased because the transmission bandwidth would be  
17      excessive and a conventional MPEG-2 decoder will not be able to handle the increased  
18      data rate or even if the decoder would have been able to support the increased data rate,  
19      such a change in the original operating conditions is not allowable. For this reason, in  
20      trick-mode, neither the original display rate of 29.97 frames per second (for NTSC or 25  
21      frames per second for PAL) nor the original transport stream (TS) multiplex rate should  
22      change. Nor is it possible to simply decimate frames since only the I frames are  
23      independently coded, and the P frames and B frames need the content of certain other

1 frames for proper decoding. The I frames typically occur once for every 15 frames.  
2 Assuming that this convention is followed in the encoding process, it would be possible  
3 to preserve and play each I frame from each and every group of pictures (GOP), resulting  
4 in a 15 times slower temporal sampling rate, or a 1 to 15 speeding up of motion if the I  
5 frames only are played back at the nominal NTSC rate of approximately 30 frames per  
6 second. Consequently, the content of a 60 minutes duration clip will be covered in 4  
7 minutes. Unfortunately the average information content per frame for the I frames is  
8 more than the average information content of I, P and B frames. Therefore, the trick-  
9 mode cannot be implemented simply by transmitting only the I frames for a speed-up by  
10 a factor of 15, because this would need an increase in the TS multiplex rate over the  
11 nominal rate.

12 In particular, in a sample analysis the average information content of an I frame  
13 has been measured to be about 56374.6 bytes. If the I frames only are transmitted at the  
14 standard NTSC rate, then the bit transmission rate would be: 8(bits per byte) \*  
15 56,374.6(bytes per frame) \* 29.97(frames per sec.) or about 13,516,374.1 bits per second  
16 only for the video stream, which is significantly above - almost 3.38 times - the original  
17 rate of 4 megabits per second used in this test. This calculation, being based on an  
18 average quantity, is ignoring the indispensable need for an actually higher transport rate  
19 to provide some safety margin to handle short-term-sustained large size I and/or P frame  
20 chains (bursts) which practically always happen. Clearly, some form of modification in  
21 the trick-mode operation definition is required to handle this problem and pull the bit-rate  
22 requirement down to the nominal 4 megabits per second.

23 Two degrees of freedom are available to achieve such a reduction in the required

1 bit-rate for trick-mode operation. The first is I frame compression quality and the second  
2 is a motion speed-up ratio. With respect to compression quality, it is well known that  
3 human observers' perception of image detail degrades with increasing motion speed of  
4 objects in the scene. Based on this fact, the type of D pictures were introduced in MPEG-  
5 1 video syntax for fast visible (forward or reverse) search purposes. (See ISO/IEC 11172-  
6 2: 1993 Information Technology - Coding of moving pictures and associated audio for  
7 digital storage media at up to about 1.5 Mbits/s - Part 2: Video, Annex D.6.6. Coding D-  
8 Pictures, p.102). D pictures make use of only the DC coefficients in intra coding to  
9 produce very low quality (in terms of SNR) reproductions of desired frames which were  
10 judged to be of adequate quality in fast search mode.

11 In order to provide support for enhanced quality trick-mode operation, the quality  
12 of the original I frames can be reduced by the preservation of just a sufficient number of  
13 AC DCT coefficients to meet the bit-rate limitation. Based on experiments with two  
14 standard video test sequences (one encoded at 15 Mbits/sec. and the other at 24  
15 Mbits/sec. and both with I frames only), it is observed that the bandwidth for I frames can  
16 be scaled to one half by keeping about 9 lowest order AC coefficients and eliminating the  
17 rest. This scheme provides good quality even at the full spatial and temporal resolution,  
18 much better than D pictures.

19 The inherent speed-up ratio lower bound imposed by the GOP structure can be  
20 relaxed and further lowered by freeze (P) frame substitution in between genuine (SNR  
21 scaled or non-scaled) I frames. The maximum number of freeze frames that can be  
22 inserted before visually disturbing motion jerkiness occurs, is very likely to depend  
23 heavily on the original GOP structure (equivalently the separation between I frames of

the original sequence) and the original amount of motion in the clip. However, 1, 2 or 3 freeze frame substitutions in between genuine I frames present reasonable choices which will yield speed-up ratios of 1 to 7.5, 1 to 5 and 1 to 3.75 respectively instead of the 1 to 15 speed-up ratio provided by the genuine I frames only implementation. (These ratios are computed by a first-order approximation that neglects a slight increase in bandwidth required by the consecutive freeze frames, which are inserted in between genuine I frames and can typically be made very small in size in comparison to the average size of a genuine I frame. Therefore, the insertion of 1, 2, 3 freeze frames will result in bandwidth reductions of 2 to 1, 3 to 1 and 4 to 1 respectively. The accuracy of this approximation degrades as more consecutive freeze frames and/or SNR scaling is employed.) An easy way to see the validity of these approximate figures is to note for example that in the case of 1 freeze frame insertion, the total presentation time of the trick-mode clip for an originally 60 minutes duration asset will increase from 4 minutes to 8 minutes. Since due to the underlying assumption of the first-order approximation stated above, the same amount of data (I frames only) will be transmitted in this doubled time interval, the bandwidth requirement will be halved. The final choice for trick-mode implementation should reflect a balanced trade-off along these two degrees of freedom. For example, SNR scaling of I frames down to 9 AC coefficients can be used along with single freeze frame insertion between I frames. These two choices, both of which are individually capable of providing a 2 to 1 bandwidth reduction as discussed before, will yield a combined 4 to 1 bandwidth reduction which will comfortably bring the non-scaled I frame-only bit-rate of 13516374.1 bits/sec. down to below the 4 Mbits/sec. quota. If the visual quality provided by 9 AC coefficients is not considered adequate, then SNR

1 scaling could be tuned to keep more AC coefficients at the expense of a smaller  
2 bandwidth reduction. This, however, could be compensated consequently by increasing  
3 the number of freeze frames to be used in between I frames. Coarser quantization (and  
4 therefore poorer visual quality) can be tolerated at high trick-mode speeds and better  
5 visual quality should be retained at lower trick-mode speeds.

6 With reference to FIG. 2, if the client has requested trick-mode operation,  
7 execution branches from step 58 to step 59. In step 59, execution branches to step 60 for  
8 a low value of speed-up. In step 60, the trick-mode stream is produced by streaming  
9 original-quality I frames and inserting three freeze frames per I frame, to yield a speed-up  
10 factor of  $15/4 = 3.75$  based on an original MPEG-2 coded stream having one I frame for  
11 every 15 frames. For a higher speed-up factor, execution branches from step 59 to step  
12 61. In step 61, either one or two freeze frames are selected per I frame to provide a  
13 speed-up factor of  $15/2 = 7.5$ , or  $15/3 = 5$  respectively. Then in step 62 the trick-mode  
14 stream is produced by streaming reduced-quality I frames and inserting the selected  
15 number of freeze frames between the reduced-quality I frames. If a trick-mode operation  
16 is not requested in step 58, then execution continues from step 58 to step 63. In step 63,  
17 the stream server computer streams original-quality MPEG-2 coded data to the client.  
18 Further details regarding trick-mode operation are described below with reference to  
19 FIGs. 7 to 10.

20

21 **II. MPEG Splicing**

22 FIGs. 3 to 6 show further details regarding use of the present invention for MPEG  
23 splicing. In particular, reduced-quality frames are substituted for the freeze frames used

1 in the seamless splicing procedure found in the common disclosure of Peter Bixby et al.,  
2 U.S. application Ser. 09/539,747 filed March 31, 2000; Daniel Gardere et al., U.S.  
3 application Ser. 09/540,347 filed March 31, 2000; and John Forecast et al. U.S.  
4 application Ser. 09/540,306 filed March 31, 2000; which are all incorporated by reference  
5 herein. The common disclosure in these U.S. applications considered pertinent to the  
6 present invention is included in the written description below with reference to FIGs. 3 to  
7 6 in the present application (which correspond to FIGs. 19, 22, 23, and 24 in each of the  
8 cited U.S. applications).

9 FIG. 3 shows a basic procedure for MPEG splicing. In the first step 121, the  
10 splicing procedure receives an indication of a desired end frame of the first clip and a  
11 desired start frame of the second clip. Next, in step 122, the splicing procedure finds the  
12 closest I frame preceding the desired start frame to be the In Point for splicing. In step  
13 123, the splicing procedure adjusts content of the first clip near the end frame of the first  
14 clip and adjusts content of the second clip near the In Point in order to reduce  
15 presentation discontinuity (due to decoder buffer underflow) and also to prevent decoder  
16 buffer overflow when decoding the spliced MPEG stream. Finally, in step 124, the  
17 concatenation of the first clip up to about the Out Point and the second clip subsequent to  
18 about the In Point is re-formatted, including re-stamping of the presentation time stamps  
19 (PTS), decoding time stamps (DTS), and program clock reference (PCR) values for the  
20 audio and video streams in the second clip.

21 Considering now video splicing, the splicing procedure should ensure the absence  
22 of objectionable video artifacts, preserve the duration of the spliced stream, and if  
23 possible, keep all of the desired frames in the spliced stream. The duration of the spliced

1 stream should be preserved in order to prevent any time drift in the scheduled play-list.  
2 In some cases, it is not possible to keep all of the original video frames due to buffer  
3 problems.

4 Management of the video buffer is an important consideration in ensuring the  
5 absence of objectionable video artifacts. In a constant bit rate (CBR) and uniform picture  
6 quality sequence, subsequent pictures typically have coded representations of drastically  
7 different sizes. The encoder must manage the decoder's buffer within several constraints.  
8 The buffer should be assumed to have a certain size defined in the MPEG-2 standard.  
9 The decoder buffer should neither overflow nor underflow. Furthermore, the decoder  
10 cannot decode a picture before it receives it in full (i.e. completely). Moreover, the  
11 decoder should not be made to "wait" for the next picture to decode; this means that  
12 every 40 ms in PAL and 1/29.97 second in NTSC, the decoder must have access to a full  
13 picture ready to be decoded.

14 The MPEG encoder manages the video decoder buffer through decode time  
15 stamps (DTS), presentation time stamps (PTS), and program clock reference (PCR)  
16 values. When splicing the end of a first clip to the beginning of a second clip, there will  
17 be a problem of video buffer management if a duration of time  $DTS_{L1}-T_e$  is different from  
18 a duration of time  $DTS_{F2}-PCR_{e2}$  minus one video frame (presentation) interval, where  
19  $DTS_{L1}$  is the DTS at the end of the first clip and indicates the time at which the video  
20 decoder buffer is emptied of video data from the first clip,  $T_e$  is the time at which the last  
21 video frame's data is finished being loaded into the video decoder buffer,  $DTS_{F2}$  is the  
22 DTS of the first frame of the second clip, and  $PCR_{e2}$  is the PCR of the second clip  
23 extrapolated from the value of the most recent received genuine PCR record, to the first

1 byte of the picture header sync word of the first video frame in the clip to start. The  
2 extrapolation adjusts this most recently received genuine PCR record value by the  
3 quotient of the displacement in data bits of the clip from the position where it appears in  
4 the second clip to the position at which video data of the first frame of the second clip  
5 begins, divided by the data transmission bit rate for transmission of the clip to the  
6 decoder. Because the time  $\text{PCR}_{e2}$  must immediately follow  $T_e$ , there will be a gap in the  
7 decoding and presentation of video frames if  $\text{DTS}_{F2}-\text{PCR}_{e2}$  is substantially greater than  
8  $\text{DTS}_{L1}-T_e$  plus one video frame interval. In this case, the buffer will not be properly full  
9 to begin decoding of the second clip one video frame interval after the last frame of the  
10 first clip has been decoded. Consequently, either the second clip will be prematurely  
11 started to be decoded or the decoder will be forced to repeat a frame one or more times  
12 after the end of the display of the last frame from the first clip to provide the required  
13 delay for the second clip's buffer build-up. In the case of a premature start for decoding  
14 the second clip, a video buffer underflow risk is generated. On the other hand, in case of  
15 repeated frames, the desired frame accuracy for scheduled play-lists is lost besides the  
16 fact that neither a guaranteed safe buffer management can be achieved through this  
17 procedure.

18 If  $\text{DTS}_{F2}-\text{PCR}_{e2}$  is substantially less than  $\text{DTS}_{L1}-T_e$  plus one video frame interval,  
19 then the decoder will not be able to decode the first frame of the second clip at the  
20 specified time  $\text{DTS}_{F2}$  because either the last frame of the first clip will not yet have been  
21 removed from the video buffer or the last frame of the first clip has already been moved  
22 but the frame interval duration required before decoding the next frame has not elapsed  
23 yet. In this case a video buffer overflow risk is generated. Video buffer overflow may

1 present a problem not only at the beginning of the second clip, but also at a subsequent  
2 location of the second clip. If the second clip is encoded by an MPEG-2 compliant  
3 encoder, then video buffer underflow or buffer overflow will not occur at any time during  
4 the decoding of the clip. However, this guarantee is no longer valid if the  $DTS_{F2}$ - $PCR_{e2}$   
5 relationship at the beginning of the second clip is altered. Consequently, to avoid buffer  
6 problems, the buffer occupancy at the end of the first clip must be modified in some  
7 fashion. This problem is inevitable when splicing between clips having significantly  
8 different ending and starting buffer levels. This is why the Society of Motion Picture and  
9 Television Engineers (SMPTE) has defined some splice types corresponding to well-  
10 defined buffer levels. (See SMPTE Standard 312M, entitled "Splice Points for MPEG-2  
11 Transport Streams," SMPTE Journal, Nov. 1998.) In order to seamlessly splice the first  
12 clip to the second clip, the content of the first clip (towards its end) is modified so that  
13  $PCR_{e2}$  can immediately follow  $T_e$  (by one byte transmission time) and  $DTS_{F2}$  can just  
14 follow  $DTS_{L1}$  (by one video frame presentation interval).

15 FIG. 4 shows a flow chart of a seamless video splicing procedure that attains the  
16 desired condition just described above. In a first step 141, the first DTS of the second  
17 clip is anchored at one frame interval later than the last DTS of the first clip in order to  
18 prevent a video decoding discontinuity. Then, in step 142, the procedure branches  
19 depending on whether the PCR extrapolated to the beginning frame of the second clip  
20 falls just after the ending time of the first clip. If so, then the splice will be seamless with  
21 respect to the original video content. Otherwise, the procedure branches to step 143. In  
22 step 143, the content of the first clip is adjusted so that the PCR extrapolated to the

1 beginning frame of the second clip falls just after the ending time of the first clip.  
2 Therefore the desired conditions for seamless video splicing are achieved.

3 With reference to FIG. 5, there is shown a more detailed flow chart of a seamless  
4 video splicing procedure. In a first step 151, the procedure inspects the content of the  
5 first clip to determine the last DTS/PTS of the first clip. This last DTS/PTS of the first  
6 clip is designated  $DTS_{L1}$ . Next, in step 152, the procedure inspects the content of the first  
7 clip to determine the time of arrival ( $T_e$ ) of the last byte of the first clip. In step 153, the  
8 procedure adds one frame interval to  $DTS_{L1}$  to find the desired first DTS location for the  
9 second clip. The sum, designated  $DTS_{F1}$ , is equal to  $DTS_{L1} + 1/FR$ , where FR is the video  
10 frame rate. In step 154, while keeping the  $DTS_{F2}$ - $PCR_{e2}$  relationship unaltered for the  
11 second clip, the procedure finds the time instant, designated  $T_s$ , at which the first byte of  
12 the second clip should arrive at the decoder buffer. This is done by calculating  
13  $T_{START}=DTS_{F2}-PCR_{e2}$ , and  $T_s=DTS_{F1}-T_{START}$ .

14 Continuing in FIG. 6, in step 155, execution branches depending on whether  $T_s$  is  
15 equal to  $T_e$  plus 8 divided by the bit rate. If not, then the clips to be spliced need  
16 modification before concatenation, and execution branches to step 156. In step 156,  
17 execution branches depending on whether  $T_s$  is less than  $T_e$  plus 8 divided by the bit rate.  
18 If not, then there is an undesired gap in between the clips to be spliced, and execution  
19 branches to step 157. In step 157, null packets are inserted into the clips to be spliced to  
20 compensate for the gap. The gap to be compensated has a number of bytes, designated  
21  $G_r$ , equal to  $(T_s-T_e)(BIT\ RATE)/8$  minus one. If in step 156,  $T_s$  is less than  $T_e$  plus 8  
22 divided by the bit rate, then execution continues from step 156 to step 158 to open up a  
23 certain amount of space in the first clip to achieve  $T_s=T_e+8/(BIT\ RATE)$ . The number of

1 bytes to drop is one plus  $(T_e - T_s)(\text{BIT RATE})/8$ . If possible, the bytes are dropped by  
2 removing null packets. Otherwise, one or more frames at the end of the first clip are  
3 replaced with corresponding reduced-quality frames, which have fewer bytes than the  
4 original-quality frames at the end of the first clip.

5 If in step 155  $T_s$  is found to be equal to  $T_e$  plus 8 divided by the bit rate, then  
6 execution continues to step 159. Execution also continues to step 159 from steps 157 and  
7 158. In step 159, the transport streams from the two clips are concatenated. Finally, in  
8 step 160, a subroutine is called to compute a video time stamp offset, designated as  
9  $V_{\text{OFFSET}}$ . This subroutine finds the DTS of the last video frame (in decode order) of the  
10 first clip. This DTS of the last video frame of the first clip is denoted  $DTS_{VL1}$ . Then the  
11 subroutine finds the original DTS of the first frame to be decoded in the second clip.  
12 This DTS of the first frame to be decoded in the second clip is denoted  $DTS_{VF2}$ . Finally,  
13 the subroutine computes the video time stamp offset  $V_{\text{OFFSET}}$  as  $DTS_{VL1} - DTS_{VF2}$  plus one  
14 video frame duration.

15

16       III. Trick Mode Operation

17 FIGs. 7 to 10 show further details regarding trick-mode operation. FIG. 7 shows  
18 a timing relationship between video presentation units (VPUs) and associated audio  
19 presentation units (APUs) in an original MPEG-2 coded data stream, and FIG. 8 shows  
20 similar timing for the fast-forward trick-mode stream produced from the original data  
21 stream of FIG. 7. (The fast-forward trick-mode stream is an example of a trick-mode  
22 stream that could be produced in step 60 of FIG. 2.) The original data stream has  
23 successive video presentation units for video frames of type I, B, B, P, B respectively.

1 The trick-mode stream has successive video presentation units for video frames of types  
2 I, F, F, I, F where "F" denotes a freeze P (or possibly B) frame. Each I frame and  
3 immediately following F frames produce the same video presentation units as a  
4 respective I frame in the original data stream of FIG. 7, and in this example, one in every  
5 15 frames in the original data stream is an I frame. Each freeze frame is coded, for  
6 example, as a P frame repeating the previous I frame or the previous P-type freeze-frame  
7 (in display order). In each freeze frame, the frame is coded as a series of maximum-size  
8 slices of macroblocks, with an initial command in each slice indicating that the first  
9 macroblock is an exact copy of the corresponding macroblock in the previous frame  
10 (achieved by predictive encoding with a zero valued forward motion compensation vector  
11 and no encoded prediction error), and two consequent commands indicating that the  
12 following macroblocks in the slice until and including the last macroblock of the slice are  
13 all coded in the same way as the first macroblock.

14 For trick-mode operation, there is also a problem of how to select audio  
15 presentation units (APU) to accompany the video presentation units that are preserved in  
16 the trick-mode stream. Because the video presentation units (VPU) have a duration of  
17 ( $1/29.97$ ) sec. or about 33.37 msec. and the audio presentation units (APU) typically have  
18 a duration of 24 msec., there is neither a one-to-one correspondence nor alignment  
19 between VPUs and APUs. In a preferred implementation, the audio content of a trick-  
20 mode clip is constructed as follows. Given the total presentation duration ( $1/29.97$ ) sec.  
21 or about 33.37 msec. for a single video frame, it is clear that always at least one and at  
22 most two 24 msec. long audio presentation units (APU) will start being presented during  
23 the end-to-end presentation interval of each video frame. This statement refers to the

1 original clip and does not consider any audio presentation unit whose presentation is  
2 possibly continuing as the video frame under consideration is just put on the display. The  
3 first of the above mentioned possibly two audio presentation units will be referred to as  
4 the aligned audio presentation unit with respect to the video frame under consideration.  
5 For example, in FIG. 8, the  $APU_j$  is the aligned audio presentation unit with respect to the  
6  $VPU_i$ . Now, when the I frames are extracted and possibly SNR scaled and possibly  
7 further interleaved with a number of freeze P frames in between them to produce the  
8 trick-mode video packetized elementary stream (PES), the associated trick-mode audio  
9 stream is constructed as follows. For each I type video frame presentation interval (and  
10 for that matter also for freeze P type video frames) in this trick-mode clip, the above  
11 stated fact of at least one (and at most two) audio presentation unit being started, holds.  
12 Then for each I frame presentation interval in the trick-mode clip, once any possibly  
13 previously started and continuing audio presentation unit ends, insert its aligned audio  
14 presentation unit (from the original clip) and continue inserting APUs from the original  
15 clip subsequent to the aligned one until covering the rest of the I frame presentation  
16 interval and also any possibly following freeze P frame presentation intervals until  
17 crossing into and overlapping (or less likely aligning) with the next I frame's presentation  
18 interval. In FIG. 8, for example, the audio presentation units  $APU_j$ ,  $APU_{j+1}$ ,  $APU_{j+2}$ , and  
19  $APU_{j+3}$  are inserted, until crossing into and overlapping with the next I frame  $VPU_{i+15}$ .  
20 Following  $APU_{j+3}$  is inserted  $APU_k$ , which designates the APU aligned with  $VPU_{i+15}$  in  
21 the original stream. Clearly, the final alignment of (the aligned and consequent) audio  
22 presentation units with respect to their associated I frames will be slightly different in the

1 trick-mode clip as compared to the original clip. However, considering how the trick-  
2 mode audio component will sound like, this poses no problem at all.

3 FIG. 9 is a flowchart of a procedure for producing the desired sequencing of audio  
4 presentation units (APUs) in the fast-forward trick-mode stream. This procedure scans  
5 the audio elementary stream in the original MPEG-2 stream to determine the sequence of  
6 APUs in the original stream and their presentation-time alignment with the I frames in the  
7 video elementary stream of the original MPEG-2 transport stream, while selecting APUs  
8 to include in the trick-mode stream. In a first step 171, execution proceeds once the end  
9 of the current APU is reached. If the end of the current APU has not entered a new VPU  
10 (*i.e.*, the beginning of the current APU is within the presentation time of one VPU and the  
11 end of the current APU is within the presentation time of the same VPU), or if it has  
12 entered a new VPU (*i.e.*, the beginning of the current APU is within the presentation time  
13 of one VPU and the end of the current APU is within the presentation time of a new  
14 (next) VPU) but the new VPU is not an I frame, then execution branches to step 174. In  
15 step 174, an APU pointer is incremented, and in step 175 execution proceeds into this  
16 next APU. If in step 173 the end of the current APU extends into an I frame, then in step  
17 176 the APU pointer is advanced to point to the first APU beginning within the duration  
18 of the VPU of the I frame in the original MPEG-2 stream.

19 FIG. 10 is a flowchart of a procedure for producing a trick-mode stream from an  
20 MPEG-2 transport stream (TS). In a first step 181, the MPEG-2 TS is inputted. In step  
21 182, the video elementary stream (VES) is extracted from the TS. In step 183, a  
22 concurrent task extracts the audio elementary stream (AES) from the TS. In step 184, I  
23 frames are extracted from the VES and valid packetized elementary stream (PES) packets

1 are formed encapsulating the I frames. In step 185, the I frames are SNR scaled, for the  
2 high speed cases of the trick-mode. In step 186, P-type freeze frames are inserted into the  
3 stream of SNR scaled I frames (in between the scaled I frames), and valid PES packets  
4 are formed for the trick-mode VES encapsulating the P-type freeze frames and the SNR  
5 scaled I frames. Concurrently, in step 187, appropriate audio access units (from the  
6 originally input MPEG-2 TS asset) are selected and concatenated based on the structure  
7 of the VES being formed for the trick-mode clip, as described above with reference to  
8 FIG. 9, and valid PES packet encapsulation is formed around these audio access units.  
9 Finally, in step 188, the trick-mode TS stream is generated by multiplexing the trick-  
10 mode VES from step 186 into a system information (SI) and audio PES carrying TS  
11 skeleton including the audio PES packets from step 187.

12

13        IV. Truncation of AC DCT Coefficients for Producing Low-Quality MPEG  
14        Coded Video

15        FIGs 11 to 19 include details of the preferred techniques for truncating AC DCT  
16 coefficients for producing low-quality MPEG-2 coded video from original-quality  
17 MPEG-2 coded video. Most of these techniques exploit the fact that in the typical  
18 (default) zig-zag scan order, the basis functions for the high-order AC DCT coefficients  
19 have an increasing frequency content. FIG 11, for example, shows a matrix of the DCT  
20 coefficients  $C_{ij}$ . The row index ( $i$ ) increases with increasing vertical spatial frequency in  
21 a corresponding 8x8 coefficient block, and the column index ( $j$ ) increases with increasing  
22 horizontal spatial frequency in the corresponding 8x8 coefficient block. The coefficient  
23  $C_{11}$  has zero frequency associated with it in both vertical and horizontal directions, and

1 therefore it is referred to as the DC coefficient of the block. The other coefficients have  
2 non-zero spatial frequencies associated with their respective basis functions, and  
3 therefore they are referred to as AC coefficients. Each coefficient has an associated basis  
4 function  $f_{ij}(x,y)$  that is separable into x and y components such that  $f_{ij}(x,y) = f_i(y)f_j(x)$ .  
5 The x and y component functions  $f_i(y)$  and  $f_j(x)$  are shown graphically in FIG. 11 as  
6 cosine functions in order to illustrate their associated spatial frequencies. In practice, the  
7 component functions are evaluated at discrete points for the 64 pixel positions in the 8x8  
8 blocks, so that each of the DCT basis functions is an 8x8 array of real numbers. In  
9 particular, the component functions are:

10

11  $f_i(y) = \text{SQRT}((2-\delta_{i-1})/8) \cos((\pi/8)(y-1/2)(i-1))$  for  $y=1, 2, 3, \dots, 8$

12  $f_j(x) = \text{SQRT}((2-\delta_{j-1})/8) \cos((\pi/8)(x-1/2)(j-1))$  for  $x=1, 2, 3, \dots, 8$

13 where  $\delta_0 = 1$ , and  $\delta_p = 0$  for  $p > 0$ . The path including a number of diagonal line segments  
14 through the matrix of coefficients in FIG. 11 denotes the default zig-zag scan order  
15 typically used for MPEG-2 encoding. Listed in this order, the coefficients are  $C_{11}, C_{12},$   
16  $C_{21}, C_{31}, C_{22}, C_{13}, C_{14}, C_{23}, C_{32}, C_{41}, \dots, C_{86}, C_{77}, C_{68}, C_{78}, C_{87}, C_{88}$ . The first coefficient  
17 in this zig-zag scan order is the DC coefficient  $C_{11}$  providing the lowest spatial frequency  
18 content in the 8x8 block of pixels, and the last coefficient in this zig-zag scan order is the  
19 coefficient  $C_{88}$  providing the highest spatial frequency content in the 8x8 block of pixels.

20 FIG. 12 is a diagram illustrating a relationship between an original MPEG-2  
21 coded bit stream 200 and a reduced-quality MPEG-2 coded bit stream 210 resulting from  
22 truncation of high-order DCT coefficients from the original MPEG-2 coded bit stream.  
23 Shown in the original MPEG-2 coded bit stream 200 is a portion of a video PES packet

1 including DCT coefficients for an 8x8 pixel block. The DCT coefficients include a  
2 differentially coded DC coefficient 201, and three (run, level) events 202, 203, 204  
3 encoding three respective nonzero AC coefficients possibly along with some zero valued  
4 AC coefficients preceding the three nonzero valued ones. The DCT coefficients are  
5 ordered according to the zig-zag scan order shown in FIG. 11 (or possibly according to an  
6 alternate zig-zag scan pattern also supported by the MPEG-2 standard), and AC  
7 coefficients having zero magnitude are described in terms of total counts of consecutive  
8 zero valued coefficients lying in between two nonzero valued coefficients, in the MPEG-  
9 2 coded bit stream. An end-of-block (EOB) code 205 signals the end of the encoded  
10 DCT coefficients for the current block. The reduced-quality MPEG-2 coded bit stream  
11 210 includes a DC coefficient 201' identical to the DC coefficient 201 in the original  
12 MPEG-2 coded bit stream 200, and a (run, level) event 202' identical to the (run, level)  
13 event 202 in the original MPEG-2 coded bit stream 200. Second and third (run, level)  
14 events, however, have been omitted from the reduced-quality MPEG-2 bit stream 210,  
15 because an EOB code 205' immediately follows the (run, level) event 202'. Therefore,  
16 the two nonzero high-order AC DCT coefficients encoded by the second and third (run,  
17 level) events 203, 204 have been omitted from the reduced-quality MPEG-2 bit stream  
18 210.

19 FIG. 13 is a flowchart of a procedure for scaling MPEG-2 coded video using a  
20 variety of techniques including the omission of AC DCT coefficients. The procedure  
21 operates upon an original-quality MPEG-2 coded video stream by removing AC DCT  
22 coefficients in this stream to produce a lower quality MPEG coded video stream. In a  
23 first step 221, execution branches to step 222 if the scaled MPEG coded video is to be

1 spatially subsampled. In step 222, the procedure removes any and all DCT coefficients  
2 for spatial frequencies in excess of the Nyquist frequency for the downsampled video.  
3 For example, if the low-quality video stream will be downsampled by a factor of two in  
4 both the vertical and the horizontal directions, then the procedure removes any and all  
5 DCT coefficients having a row index (i) greater than four and any and all DCT  
6 coefficients having a column index (j) greater than four. This requires the decoding of  
7 the (run, level) coded coefficients to the extent necessary to obtain an indication of the  
8 coefficient indices. If a sufficient number of the original AC DCT coefficients are  
9 removed for a desired bandwidth reduction, then the scaling procedure is finished.  
10 Otherwise, execution branches from step 223 to step 224. Execution also continues from  
11 step 221 to step 224 if spatial downsampling is not intended.

12 In step 224, execution branches to step 225 if low-pass scaling is desired. Low-  
13 pass scaling requires the least computational resources and may produce the best results  
14 if the scaled, low-quality MPEG coded video is spatially downsampled. In step 225, the  
15 procedure retains up to a certain number of lowest-order AC DCT coefficients for each  
16 block and removes any additional DCT coefficients for each block. This is a kind of  
17 frequency domain signal-to-noise ratio scaling (FDSNR) that will be designated  
18 FDSNR\_LP. A specific example of the procedure for step 225 will be described below  
19 with reference to FIG. 14.

20 Execution continues from step 224 to step 226 if low-pass scaling is not desired.  
21 In step 226, execution branches to step 227 if largest magnitude based scaling is desired.  
22 Largest magnitude based scaling produces the least squared error or difference between  
23 the original-quality MPEG-2 coded video and the reduced-quality MPEG coded video for

1 a given number of nonzero AC coefficients to preserve, but it requires more  
2 computational resources than the low-pass scaling of step 225. More computational  
3 resources are needed because if there are more nonzero AC coefficients than the desired  
4 number of AC coefficients for a block, then the (run, level) events must be decoded fully  
5 to obtain the coefficient magnitudes, and additional resources are required to find the  
6 largest magnitude coefficients. In step 227, the procedure retains up to a certain number  
7 of largest magnitude AC DCT coefficients for each block, and removes any and all  
8 additional AC DCT coefficients for each block. This is a kind of frequency domain  
9 signal-to-noise ratio scaling (FDSNR) that will be designated FDSNR\_LM. A specific  
10 example of the procedure for step 227 will be described below with reference to FIG. 15.

11 If in step 226 largest magnitude based scaling is not desired, then execution  
12 continues to step 228. In step 228, execution branches to step 229 to retain up to a certain  
13 number of AC DCT coefficients that differ in magnitude from up to that number of  
14 largest magnitude AC DCT coefficients by no more than a certain limit. This permits a  
15 kind of approximation to FDSNR\_LM in which an approximate search is undertaken for  
16 the largest magnitude AC DCT coefficients if there are more nonzero AC DCT  
17 coefficients than the desired number of AC DCT coefficients in a block. The  
18 approximate search can be undertaken using a coefficient magnitude classification  
19 technique such as a hashing technique, and the low-pass scaling technique can be applied  
20 to the classification level that is incapable of discriminating between the desired number  
21 of largest magnitude AC DCT coefficients. A specific example is described below with  
22 reference to FIG. 19.

1       With reference to FIG. 14, there is shown a flowchart of a procedure for scaling  
2       MPEG-2 coded video using the low-pass frequency-domain signal-to-noise (FDSNR\_LP)  
3       scaling technique. This procedure scans and selectively copies components of an input  
4       stream of original-quality MPEG-2 coded video to produce an output stream of reduced-  
5       quality MPEG-2 coded video. The procedure is successively called, and each call  
6       processes coefficient data in the input stream for one 8x8 block of pixels. No more than a  
7       selected number “k” of coded lowest order (nonzero or zero valued) AC coefficients are  
8       copied for the block where the parameter “k” can be specified for each block.

9           In a first step 241 of FIG. 14, the procedure parses and copies the stream of  
10       original-quality MPEG-2 coded data up to and including the differential DC coefficient  
11       variable-length code (VLC). Next, in step 242, a counter variable “l” is set to zero. In  
12       step 243, the procedure parses the next (run, level) event VLC in the stream of original-  
13       quality MPEG-2 coded data. In step 244, if the VLC just parsed is an end-of-block  
14       (EOB) marker, execution branches to step 245 to copy the VLC to the stream of reduced-  
15       quality MPEG-2 coded video, and the procedure is finished for the current block.

16           In step 244, if the VLC just parsed is not an EOB marker, then execution  
17       continues to step 246. In step 246, a variable “r” is set equal to the run length of zeroes  
18       for the current (run, level) event, in order to compute a new counter value  $l+r+1$ . In step  
19       247, if the new counter value  $l+r+1$  is greater than the parameter “k”, then the procedure  
20       branches to step 248 to copy an EOB marker to the stream of reduced-quality MPEG  
21       coded data. After step 248, execution continues to step 249, where the procedure parses  
22       the input stream of original-quality MPEG-2 coded data until the end of the first EOB  
23       marker, and the procedure is finished for the current block.

1        In step 247, if the new counter value  $l+r+1$  is not greater than the parameter “k”,  
2        then execution continues to step 250. In step 250, execution branches to step 251 if the  
3        new counter value  $l+r+1$  is not equal to “k” (which would be the case if the new counter  
4        value is less than “k”). In step 251, the counter state  $l$  is set equal to the new counter  
5        value  $l+r+1$ . Then, in step 252, the VLC just parsed (which will be a VLC encoding a  
6        (run, level) event) is copied from the stream of original-quality MPEG-2 coded data to  
7        the stream of reduced-quality MPEG-2 coded data. After step 252, execution loops back  
8        to step 243 to continue the scanning of the stream of original-quality MPEG-2 coded  
9        data.

10      In step 250, if the new counter value  $l+r+1$  is equal to “k”, then execution  
11     branches from step 250 to step 253, to copy the VLC just parsed (which will be a VLC  
12     encoding a (run, level) event) from the stream of original-quality MPEG-2 coded data to  
13     the stream of reduced-quality MPEG-2 coded data. Next, in step 254, the procedure  
14     copies an EOB marker to the stream of reduced-quality MPEG-2 coded data. After step  
15     254, execution continues to step 249, where the procedure parses the input stream of  
16     original-quality MPEG-2 coded data until the end of the first EOB marker, and the  
17     procedure is finished for the current block.

18      FIG. 15 is a flowchart of a procedure for scaling MPEG-2 coded video using the  
19     largest magnitude based frequency-domain signal-to-noise ratio (FDSNR\_LM) scaling  
20     technique. This routine is successively called, and each call processes coefficient data in  
21     the input stream for one 8x8 block of pixels. No more than a specified number “k” of  
22     largest magnitude AC DCT coefficients are copied for the block, and a different number  
23     “k” can be specified for each block.

1        In a first step 261 in FIG. 15, the procedure parses and copies the input stream of  
2 original-quality MPEG-2 coded data to the output stream of lower-quality MPEG-2 data  
3 up to and including the differential DC coefficient variable-length code (VLC). Then in  
4 step 262 all (run, level) event VLCs are parsed and decoded until and including the EOB  
5 marker of the current block. The decoding produces coefficient identifiers and  
6 corresponding quantization indices representing the quantized coefficient values. In step  
7 263, the quantization indices are transformed to quantized coefficient values. In step 264,  
8 the (quantized) coefficients are sorted in descending order of their magnitudes. In step  
9 265, the first “k” coefficients of the sorted list are preserved and the last 63-k AC DCT  
10 coefficients in the sorted list are set to zero. In step 266, (run, level) event formation and  
11 entropy coding (VLC encoding) are applied to the new set of coefficient values. Finally,  
12 in step 267, the VLCs resulting from step 266 are copied to the output stream until and  
13 including the EOB marker.

14        The sorting step 264 of the FDSNR\_LM procedure can consume considerable  
15 computational resources. It is important to notice that not a full sorting of the quantized  
16 AC coefficients with respect to their magnitudes but rather a search for a specified  
17 number “k” of largest magnitude AC coefficients is all that is required. This task can be  
18 performed exactly or approximately in different ways so as to avoid the complexity  
19 associated with a conventional sorting procedure. In general, a relatively large number of  
20 the 63 AC DCT coefficients will have a quantized value of zero. Only the non-zero  
21 coefficients need be included in the sorting process. Moreover, if there are “n” non-zero  
22 coefficients and only “k” of them having the largest magnitudes are to be preserved in the  
23 output stream, then the sorting process may be terminated immediately after only the

1 largest magnitude “k” coefficients have been found, or equivalently immediately after  
2 only the smallest magnitude “n-k” coefficients have been found. Moreover, the sorting  
3 procedure itself can be different depending on a comparison of “k” to “n” in order to  
4 minimize computations.

5 With reference to FIG. 16, there is shown a flowchart of a procedure that selects  
6 one of a number of techniques for finding a certain number “k” of largest values out of a  
7 set of “n” values. In a first step 271, execution branches to step 272 if “k” is less than  $\frac{1}{2}$   
8 “n.” In step 272, execution branches to step 273 if “k” is much less than  $\frac{1}{2}$  “n.” In step  
9 273, the first “k” values are sorted to produce a list of “k” sorted values, and then the last  
10 “n-k” values are scanned for any value greater than the minimum of the sorted “k”  
11 values. If a value greater than the minimum of the sorted “k” values is found, then that  
12 minimum value is removed and the value greater than the minimum value is inserted into  
13 the list of “k” sorted values. At the end of this procedure, the list of sorted “k” values  
14 will contain the maximum “k” values out of the original “n” values. A specific example  
15 of this procedure is described below with reference to FIG. 17.

16 In step 272, if “k” is not much less than  $\frac{1}{2}$  “n”, then execution branches to step  
17 274. In step 274, a bubble-sort procedure is used, including “k” bottom-up bubble-sort  
18 passes over the “n” values to put “k” maximum values on top of a sorting table. An  
19 example of such a bubble-sort procedure is listed below:

20

```
21 /* TABLE(0) to TABLE(n-1) INCLUDES n VALUES */  
22 /* MOVE THE k LARGEST OF THE n VALUES IN TABLE TO THE RANGE  
23 TABLE(0) TO TABLE(k-1) IN THE TABLE */
```

```
1    /* k <= ½ n */  
2    FOR i=1 to k  
3        FOR j=1 to n-i  
4            IF (TABLE(n-j) > TABLE(n-j-1)) THEN(  
5                /* SWAP TABLE(n-j) WITH TABLE(n-j-1) */  
6                TEMP ← TABLE(n-j)  
7                TABLE(n-j) ← TABLE(n-j-1)  
8                TABLE(n-j-1) ← TEMP)  
9            NEXT j  
10           NEXT i  
11
```

12        In step 271, if “k” is not less than  $\frac{1}{2}$  “n”, then execution branches to step 275. In  
13 step 275, if “k” is much greater than  $\frac{1}{2}$  “n”, then execution branches to step 276. In step  
14 276, a procedure similar to step 273 is used, except the “n-k” minimum values are  
15 maintained in a sorted list, instead of the “k” maximum values. In step 276, the last “n-k”  
16 values are placed in the sort list and sorted, and then the first “k” values are scanned for  
17 any value less than the maximum value in the sorted list. If a value less than the  
18 maximum value in the sorted list is found, then the maximum value in the sorted list is  
19 removed, and the value less than this maximum value is inserted into the sorted list. At  
20 the end of this procedure, the values in the sorted list are the “n-k” smallest values, and  
21 the “k” values excluded from the sorted list are the “k” largest values.

22        In step 275, if “k” is not much greater than  $\frac{1}{2}$  “n”, then execution branches to step  
23 277. In step 277, a bubble-sort procedure is used, including “n-k” top-down bubble-sort

1 passes over the "n" values to put "n-k" minimum values at the bottom of a sorting table.  
2 Consequently, the k maximum values will appear in the top "k" entries of the table. An  
3 example of such a bubble-sort procedure is listed below:

4

```
5 /* TABLE(0) to TABLE(n-1) INCLUDES n VALUES */  
6 /* MOVE THE n-k SMALLEST OF THE n VALUES IN THE TABLE */  
7 /* TO THE RANGE TABLE(k) TO TABLE(n-1) IN THE TABLE */  
8 /* n > k >= 1/2 n */  
9 FOR i=1 to n-k  
10 FOR j=0 to n-i-1  
11 IF (TABLE(j) < TABLE(j+1)) THEN(  
12     /* SWAP TABLE(j) WITH TABLE(j+1)*/  
13     TEMP ← TABLE(j)  
14     TABLE(j) ← TABLE(j+1)  
15     TABLE(j+1) ← TEMP)  
16     NEXT j  
17     NEXT i  
18
```

19 Turning now to FIG. 17, there is shown a flowchart of a procedure for finding up  
20 to a specified number "k" of largest magnitude AC DCT coefficients from a set of "n"  
21 coefficients, corresponding to the procedure of FIG. 16 for the case of  $k << \frac{1}{2}n$ . In a first  
22 step 281, a counter "i" is set to zero. In step 282, the next AC DCT coefficient is  
23 obtained from the input stream of original-quality MPEG-2 coded data. If an EOB

1 marker is reached, as tested in step 283, then execution returns. In step 284, the counter  
2 “i” is compared to the specified number “k”, and if “i” is less than “k”, execution  
3 continues to step 285. In step 285, a coefficient index and magnitude for the AC DCT  
4 coefficient is placed on a sort list. In step 286, the counter “i” is incremented, and  
5 execution loops back to step 282.

6 Once the sort list has been loaded with indices and magnitudes for “k” AC DCT  
7 coefficients and one additional coefficient has been obtained from the input stream,  
8 execution branches from step 284 to step 287. In step 287 the list is sorted by magnitude,  
9 so that the minimum magnitude appears at the end of the list. Then in step 288 the  
10 coefficient magnitude of the current coefficient last obtained from the input stream is  
11 compared to the magnitude at the end of the list. If the coefficient magnitude of the  
12 current coefficient is not greater than the magnitude appearing at the end of the list, then  
13 execution continues to step 289 to get the next AC DCT coefficient from the input  
14 stream. If an EOB marker is reached, as tested in step 290, then execution returns.  
15 Otherwise, execution loops back to step 288.

16 In step 288, if the magnitude of the current coefficient is greater than the  
17 magnitude at the end of the list, then execution branches to step 291. In step 291, the  
18 entry at the end of the list is removed. In step 292, a binary search is performed to  
19 determine the rank position of the magnitude of the current coefficient, and in step 293,  
20 the current coefficient index and magnitude are inserted into the list at the rank position.  
21 The list, for example, is a linked list in the conventional fashion to facilitate the insertion  
22 of an entry for the current coefficient at any position in the list. After step 293, execution  
23 loops back to step 288.

1 An approximation technique of coefficient magnitude classification can be used to  
2 reduce the computational burden of sorting by coefficient magnitude. A specific example  
3 is the use of hashing of the coefficient magnitude and maintaining lists of the indices of  
4 coefficients having the same magnitude classifications. As shown in FIG. 18, a hash  
5 table 300 is linked to hash lists 301 storing the indices of classified coefficients. As  
6 shown, the hash table 300 is a list of  $2^M$  entries, where "M" is three, and an entry has a  
7 value of zero if its associated list is empty, and otherwise the entry has a pointer to the  
8 end of the coefficients in its associated list. The lists shown in FIG. 18 have fixed  
9 memory allocations in which the pointers in the hash table also indicate the number of  
10 coefficient indices in the respective hash lists. Alternatively, the hash lists could be  
11 dynamically allocated and linked in the conventional fashion.

12 FIG. 19 shows a flowchart of a procedure for using the hash table 300 and hash  
13 lists 301 of FIG. 18 to perform a sort of "k" coefficients having approximately the largest  
14 magnitudes from a set of "n" coefficients. This approximation technique ensures that  
15 none of the "k" coefficients selected will have a magnitude that differs by more than a  
16 certain error limit from the smallest magnitude value of "k" coefficients having the  
17 largest magnitude. The error limit is established by the number of hash table entries, and  
18 it is the range of the magnitudes that can be hashed to the same hash table entry.

19 In a first step 311 in FIG. 19, the hash table is cleared. Then in step 312, the next  
20 AC DCT coefficient is obtained from the input stream. If an EOB marker is not reached,  
21 as tested in step 313, then execution continues to step 314. In step 314, a hash table  
22 index is stripped from the most significant bits (MSBs) of the coefficient magnitude. For  
23 the hash table in FIG. 18 having eight entries, the three most significant bits of the

1 coefficient magnitude are stripped from the coefficient magnitude. This is done by a bit  
2 masking operation together with a logical arithmetic shift operation. Then in step 315,  
3 the coefficient index is inserted on the hash list of the indexed hash table entry. For  
4 example, the hash table entry is indexed to find the pointer to where the coefficient index  
5 should be inserted, and then the pointer in the hash table entry is incremented. After step  
6 315, execution loops back to step 312. Once all of the AC coefficients for the block have  
7 been classified by inserting them in the appropriate hash lists, an EOB marker will be  
8 reached, and execution will branch from step 313 to step 316.

9 Beginning in step 316, the hash table and hash lists are scanned to find  
10 approximately the “k” largest magnitude coefficients. The hash lists linked to the bottom  
11 entries of the hash table will have the indices for the largest magnitude coefficients. Each  
12 hash list is scanned from its first entry to its last entry, so that each hash list is accessed as  
13 a first-in-first-out queue. Therefore, in each magnitude classification, the coefficient  
14 ordering in the output stream will be the same as the coefficient ordering in the input  
15 stream, and the approximation will have a “low pass” effect in which possibly some  
16 lower-frequency coefficients having slightly smaller magnitudes will be retained at the  
17 expense of discarding some higher-frequency coefficients having slightly larger  
18 magnitudes. (The approximation results from the fact that the last hash list to be scanned  
19 is not itself sorted, and to eliminate the error of the approximation, the last hash list to be  
20 scanned could be sorted.)

21 In step 316, a scan index “i” is set to  $2^M-1$  in order to index the hash table  
22 beginning at the bottom of the table, and a counter “j” is set equal to “k” in order to stop  
23 the scanning process after finding “k” coefficients. Next, in step 317, the hash table is

1 indexed with “i”. In step 318, if the indexed entry of the hash table is zero, then  
2 execution branches to step 319. In step 319, the procedure is finished if “i” is equal to  
3 zero; otherwise, execution continues to step 320. In step 320, the index “i” is  
4 decremented, and execution loops back to step 317.

5 If in step 318 the indexed hash table entry is not zero, then execution continues to  
6 step 321. In step 321, the next entry is obtained from the indexed hash list, and the  
7 coefficient index in the entry is used to put the indexed coefficient in the output stream.  
8 Then in step 322 the counter “j” is decremented, and in step 323 the counter “j” is  
9 compared to zero. In step 323, if the counter “j” is less than or equal to zero, then the  
10 procedure is finished. Otherwise, if the counter “j” is not less than or equal to zero in  
11 step 323, execution branches to step 324. In step 324, if the end of the hash list has not  
12 been reached, execution loops back to step 321 to get the next entry in the hash list.  
13 Otherwise, if the end of the hash list has been reached, execution branches to step 319.

14 The FDSNR\_LM procedure, as described above, in general provides a significant  
15 improvement in peak signal-to-noise ratio (PSNR) over the FDSNR\_LP procedure when  
16 each procedure retains the same number of non-zero AC DCT coefficients. It has been  
17 found, however, that substantially more bits are required for the (run, level) coding of the  
18 non-zero AC DCT coefficients resulting from the FDSNR\_LM procedure than those  
19 resulting from the FDSNR\_LP procedure, provided that the same coefficient quantization  
20 and scanning method is used. Therefore, the FDSNR\_LM procedure provides at best a  
21 marginal improvement in rate-distortion (PSNR as a function of bit rate) over the  
22 FDSNR\_LP procedure unless the non-zero AC DCT coefficients for the FDSNR\_LM  
23 procedure are quantized, scanned, and/or (run, level) coded in a fashion different from the

1 quantization, scanning, and/or (run, level) coding of the coefficients in the original  
2 MPEG-2 clip. A study of this problem resulted in a discovery that it is sometimes  
3 possible to reduce the number of bits for (run, level) coding of coefficients for an 8x8  
4 block including a given number of the non-zero largest magnitude AC DCT coefficients  
5 if additional coefficients are also (run, level) coded for the block.

6 The (run, level) coding of the non-zero AC DCT coefficients from the  
7 FDSNR\_LM procedure has been found to require more bits than from the FDSNR\_LP  
8 procedure due to an increased occurrence frequency of escape sequences for the (run,  
9 level) coding. The increased frequency of escape sequences is an indication that the  
10 statistical likelihood of possible (run, level) combinations for the non-zero AC DCT  
11 coefficients selected by the FDSNR\_LM procedure is different from the statistical  
12 likelihood of possible (run, level) combinations for the non-zero AC DCT coefficients  
13 produced by the standard MPEG-2 coding process and in particular those selected by the  
14 FDSNR\_LP procedure.

15 The MPEG-2 coding scheme assigns special symbols to the (run, level)  
16 combinations that occur very frequently in ordinary MPEG-2 coded video. The most  
17 frequent (run, level) combinations occur for short run lengths (within the range of about  
18 0 to 5, where the run length can range from 0 to 63) and relatively low levels (about 1 to  
19 10, where the level can range from 1 to 2048). The most frequent of these special  
20 symbols are assigned variable-length code words (VLCs). If a (run, level) combination  
21 does not have such a VLC, then it is coded with an escape sequence composed of a 6-bit  
22 escape sequence header code word followed by a 6-bit run length followed by a 12 bit  
23 signed level. An escape sequence requires a much greater number of bits than the VLCs

1 which have varying lengths depending on their relative frequency. In particular, the  
2 escape sequences each has 24 bits, and the variable-length code words have a maximum  
3 of 17 bits.

4 There are two (run, level) VLC tables in MPEG-2. The first coding table is  
5 designated TABLE 0, and the second is designated as TABLE 1. These tables specify  
6 the (run, level) combinations having VLCs. For each table, the (run, level) combinations  
7 represented by VLCs and the range of the VLC lengths are summarized below:

8

9 SUMMARY OF PROPERTIES OF DCT COEFFICIENT TABLE ZERO

10 (Table Zero is Table B.14, p. 135 of ISO/IEC 13818-2 1996E)

11

12	<u>Run</u>	<u>Range of Levels</u>	<u>Range of Code Lengths</u>
13	0	1 to 40	2 to 16
14	1	1 to 18	4 to 17
15	2	1 to 5	5 to 14
16	3	1 to 4	6 to 14
17	4	1 to 3	6 to 13
18	5	1 to 3	7 to 14
19	6	1 to 3	7 to 17
20	7	1 to 2	7 to 13
21	8	1 to 2	8 to 13
22	9	1 to 2	8 to 14
23	10	1 to 2	9 to 14

1	11	1 to 2	9 to 17
2	12	1 to 2	9 to 17
3	13	1 to 2	9 to 17
4	14	1 to 2	11 to 17
5	15	1 to 2	11 to 17
6	16	1 to 2	11 to 17
7	17	1	13
8	18	1	13
9	19	1	13
10	20	1	13
11	21	1	13
12	22	1	14
13	23	1	14
14	24	1	14
15	25	1	14
16	26	1	14
17	27	1	17
18	28	1	17
19	29	1	17
20	30	1	17
21	31	1	17
22			

23 SUMMARY OF PROPERTIES OF DCT COEFFICIENT TABLE ONE

1 (Table One is Table B.15, p. 139 of ISO/IEC 13818-2 1996E)

2

3	<u>Run</u>	<u>Range of Levels</u>	<u>Range of Code Lengths</u>
4	0	1 to 40	3 to 16
5	1	1 to 18	4 to 17
6	2	1 to 5	6 to 14
7	3	1 to 4	6 to 14
8	4	1 to 3	7 to 13
9	5	1 to 3	7 to 14
10	6	1 to 3	8 to 17
11	7	1 to 2	8 to 13
12	8	1 to 2	8 to 13
13	9	1 to 2	8 to 14
14	10	1 to 2	8 to 14
15	11	1 to 2	9 to 17
16	12	1 to 2	9 to 17
17	13	1 to 2	9 to 17
18	14	1 to 2	10 to 17
19	15	1 to 2	10 to 17
20	16	1 to 2	11 to 17
21	17	1	13
22	18	1	13
23	19	1	13

1	20	1	13
2	21	1	13
3	22	1	14
4	23	1	14
5	24	1	14
6	25	1	14
7	26	1	14
8	27	1	17
9	28	1	17
10	29	1	17
11	30	1	17
12	31	1	17
13			

The FDSNR\_LP procedure selected AC DCT coefficients have (run, level) symbol statistics that are similar to the statistics of ordinary MPEG-2 coded video, and therefore the FDSNR\_LP AC DCT coefficients have a similar frequency of occurrence for escape sequences in comparison to the ordinary MPEG-2 coded video. In contrast, the FDSNR\_LM procedure selects AC DCT coefficients resulting in (run, level) combinations that are less likely to be encountered in ordinary MPEG-2 coded video. This is due to two reasons. First, the FDSNR\_LM procedure selects AC DCT coefficients having the largest levels. Second, the FDSNR\_LM procedure introduces longer run lengths due to the elimination of coefficients over the entire range of coefficient indices. The result is a significantly increased rate of occurrence for escape

1 sequences. Escape sequences form the most inefficient mode of coefficient information  
2 encoding in MPEG-2 incorporated into the standard so as to cover important but very  
3 rarely occurring coefficient information.

4 In order to improve the rate-distortion performance of the scaled-quality MPEG-2  
5 coded video resulting from the FDSNR\_LM procedure, the non-zero AC DCT  
6 coefficients selected by the FDSNR\_LM procedure should be quantized, scanned, and/or  
7 (run, level) coded in such a way that tends to reduce the frequency of the escape  
8 sequences. For example, if the original-quality MPEG-2 coded video was (run, level)  
9 coded using TABLE 0, then the largest magnitude coefficients should be re-coded using  
10 TABLE 1 because TABLE 1 provides shorter length VLCs for some (run, level)  
11 combinations having higher run lengths and higher levels. It is also possible that re-  
12 coding using the alternate scan method instead of the zig-zag scan method may result in a  
13 lower frequency of occurrence for escape sequences. For example, each picture could be  
14 (run, level) coded for both zig-zag scanning and alternate scanning, and the scanning  
15 method providing the fewest escape sequences, or the least number of bits total, could be  
16 selected for the coding of the reduced-quality coded MPEG video.

17 There are two methods having general applicability for reducing the frequency of  
18 escape sequences resulting from the FDSNR\_LM procedure. The first method is to  
19 introduce a non-zero, “non-qualifying” AC DCT coefficient of the 8x8 block into the list  
20 of non-zero qualifying AC DCT coefficients to be coded for the block. In this context, a  
21 “qualifying” coefficient is one of the k largest magnitude coefficients selected by the  
22 FDSNR\_LM procedure. The non-qualifying coefficient referred to above, must be lying  
23 in between two qualifying AC DCT coefficients (in the coefficient scanning order) that

1 generate the (run, level) combination causing the escape sequence. Moreover, this non-  
2 qualifying coefficient must cause the escape sequence to be replaced with two shorter  
3 length VLCs when the AC DCT coefficients are (run, level) coded. This first method has  
4 the effect of not only decreasing the number of bits in the coded reduced-quality MPEG  
5 video in most cases, but also increasing the PSNR.

6 The qualifying AC DCT coefficient causing the escape sequence that is first in the  
7 coefficient scanning order will be simply referred to as the first qualifying coefficient.  
8 The qualifying AC DCT coefficient causing the escape sequence that is second in the  
9 coefficient scanning order will be simply referred to as the second qualifying coefficient.  
10 For example, suppose the qualifying coefficients in zig-zag scan order for an 8x8 block  
11 include C<sub>51</sub> followed by C<sub>15</sub> having a level of 40. If only the qualifying coefficients were  
12 (run, level) coded for the microblock, C<sub>15</sub> would result in a run length of 3, because there  
13 are a total of three non-qualifying coefficients (C<sub>42</sub>, C<sub>33</sub>, and C<sub>24</sub>) between C<sub>51</sub> and C<sub>15</sub> in  
14 the scan order. Therefore, C<sub>15</sub> would have to be coded as an escape sequence, because a  
15 run of 3 and level of 40 does not have a special symbol. In this example, the escape  
16 sequence is in effect caused by a first qualifying coefficient, which is C<sub>51</sub>, and a second  
17 qualifying coefficient, which is C<sub>15</sub>. This escape sequence can possibly be eliminated  
18 say, if C<sub>24</sub> is a non-zero, non-qualifying coefficient of the block, C<sub>24</sub> has a level of 5 or  
19 less, and C<sub>24</sub> is (run, level) coded together with the qualifying coefficients. For example,  
20 assuming that C<sub>24</sub> has a level of 5, and using the MPEG-2 (run, level) coding TABLE 1,  
21 then C<sub>24</sub> has a run length of two and is coded as the special symbol 0000 0000 1010 0s,  
22 where "s" is a sign bit, and C<sub>15</sub> now has a run length of 0 and is coded as the special  
23 symbol 0000 0000 0010 00s. Such a consideration clearly applies to the rest of the non-

1 zero non-qualifying coefficients lying in between the two qualifying coefficients  
2 producing the escape sequence. In the above example, these non-qualifying coefficients  
3 are  $C_{42}$  and  $C_{33}$ .

4 Whether or not an escape sequence can be eliminated from the (run, level) coding  
5 of the qualifying coefficients can be determined by testing a sequence of conditions. The  
6 first condition is that the second qualifying coefficient must have a level that is not  
7 greater than the maximum level of 40 for the special (run, level) symbols. If this  
8 condition is satisfied, then there must be a non-zero non-qualifying AC DCT coefficient  
9 that is between the first and second qualifying coefficients in the coefficient scanning  
10 order. If there is such a non-qualifying coefficient, then the combination of its level and  
11 the run length between the first qualifying coefficient and itself in the coefficient  
12 scanning order must be one of the special (run, level) symbols. If so, then the  
13 combination of the level of the second qualifying coefficient and the run length between  
14 the non-qualifying coefficient and the second qualifying coefficient must also be a special  
15 (run, level) symbol, and if so, all required conditions have been satisfied. If not, then the  
16 conditions with respect to the non-qualifying coefficient are successively applied to any  
17 other non-zero non-qualifying AC DCT coefficient of the block lying in between the two  
18 qualifying coefficients, until either all conditions are found to be satisfied or all such non-  
19 qualifying coefficients are tested and failed. If there are sufficient computational  
20 resources, this search procedure should be continued to find all such non-qualifying  
21 coefficients that would eliminate the escape sequence, and to select the non-qualifying  
22 coefficient that converts the escape sequence to the pair of special symbols having  
23 respective code words that in combination have the shortest length.

1 A flow chart for a modified FDSNR\_LM procedure using the first method is  
2 shown in FIGS. 20 and 21. In a first step 331 of FIG. 20, the procedure finds up to "k"  
3 largest magnitude non-zero AC DCT coefficients (i.e., the "qualifying coefficients") for  
4 the block. (This first step 331 is similar to steps 261 to 265 of FIG. 15, as described  
5 above.) In step 332, (run, level) coding of the qualifying coefficients is begun in the scan  
6 order using the second coding table (Table 1). This (run, level) coding continues until an  
7 escape sequence is reached in step 333, or the end of the block is reached in step 336. If  
8 an escape sequence is reached, execution branches from step 333 to step 334. If the level  
9 of the second qualifying coefficient causing the escape sequence is greater than 40,  
10 execution continues from step 334 to step 336. Otherwise, execution branches from step  
11 334 to step 335 to invoke a subroutine (as further described below with reference to FIG.  
12 21) to possibly include a non-zero non-qualifying AC DCT coefficient in the (run, level)  
13 coding to eliminate the escape sequence. The subroutine either returns without success,  
14 or returns such a non-qualifying coefficient so that the escape sequence is replaced with  
15 the two new (run, level) codings of the first qualifying coefficient and the non-qualifying  
16 coefficient and then the non-qualifying coefficient and the second qualifying coefficient.  
17 From step 335, execution continues to step 336. Execution returns from step 336 if the  
18 end of the block is reached. Otherwise, execution continues from step 336 to step 337, to  
19 continue (run, level) coding of the qualifying coefficients in the scan order using the  
20 second coding table (TABLE 1). This (run, level) coding continues until an escape  
21 sequence results, as tested in step 333, or until the end of the block is reached, as tested in  
22 step 336.

With reference to FIG. 21, there is shown a flow chart of the subroutine (that was called in step 335 of FIG. 20) for attempting to find a non-zero, non-qualifying AC DCT coefficient that can be (run, level) coded to eliminate an escape sequence for a qualifying coefficient. In a first step 341, the procedure identifies the first qualifying coefficient and the second qualifying coefficient causing the escape sequence. For example, the subroutine of FIG. 21 can be programmed as a function having, as parameters, a pointer to a list of the non-zero AC DCT coefficients in the scan order, an index to the first qualifying coefficient in the list, and an index to the second qualifying coefficient in the list. In step 342, the subroutine looks for a non-zero non-qualifying AC DCT coefficient between the first and the second qualifying coefficients in the scan order. For example, the value of the index to the first qualifying coefficient is incremented and compared to the value of the index for the second qualifying coefficient, and if they are the same, there is no such non-qualifying coefficient. Otherwise, if the new coefficient pointed to (by incrementing the index of the first qualifying coefficient) is a non-zero coefficient then it becomes a candidate non-qualifying coefficient deserving further testing. If however the new coefficient pointed to (by incrementing the index of the first qualifying coefficient) has a value zero then it is not a candidate non-qualifying coefficient. If no such (candidate) non-qualifying coefficients are found, as tested in step 343, then execution returns from the subroutine with a return code indicating that the search has been unsuccessful. Otherwise, execution continues to step 344.

In step 344, the non-qualifying coefficient is (run, level) coded, to determine in step 345 whether it codes to an escape sequence. If it codes to an escape sequence, then execution loops back from step 345 to step 342 to look for another non-zero non-

1 qualifying AC DCT coefficient in the scan order between the first and second qualifying  
2 coefficients. If it does not code to an escape sequence, then execution continues from  
3 step 345 to step 346. In step 346, the second qualifying coefficient is (run, level) coded,  
4 using the new run length, which is the number of coefficients in the scan order between  
5 the non-qualifying coefficient and the second qualifying coefficient. If it codes to an  
6 escape sequence, as tested in step 347, then execution loops back from step 347 to step  
7 342 to look for another non-zero non-qualifying AC DCT coefficient in the scan order  
8 between the first and second qualifying coefficients. If it does not code to an escape  
9 sequence, then execution continues from step 347 to step 348.

10 In step 348, execution returns with a successful search result unless a continue  
11 search option has been selected. If the continue search option has been selected, then  
12 execution branches from step 348 to step 349 to search for additional non-zero non-  
13 qualifying AC DCT coefficients that would eliminate the escape sequence. In other  
14 words, steps 342 to 347 are repeated in an attempt to find additional non-zero non-  
15 qualifying AC DCT coefficients that would eliminate the escape sequence. If no more  
16 such non-qualifying coefficients are found, as tested in step 350, execution returns with a  
17 successful search result. Otherwise, execution branches from step 350 to step 351 to  
18 select the non-qualifying coefficient giving the shortest overall code word length and/or  
19 the largest magnitude for the best PSNR, and execution returns with a successful search  
20 result. For example, for each non-qualifying coefficient that would eliminate the escape  
21 sequence, the total bit count is computed for the (run, level) coding of the non-qualifying  
22 coefficient and the second qualifying coefficient. Then a search is made for the non-  
23 qualifying coefficient producing the smallest total bit count, and if two non-qualifying

1 coefficients which produce the same total bit count are found, then the one having the  
2 largest level is selected for the elimination of the escape sequence.

3 A second method of reducing the frequency of occurrence of the escape  
4 sequences in the (run, level) coding of largest magnitude AC DCT coefficients for an 8x8  
5 block is to change the mapping of coefficient magnitudes to the levels so as to reduce the  
6 levels. Reduction of the levels increases the likelihood that the (run, level) combinations  
7 will have special symbols and therefore will not generate escape sequences. This second  
8 method has the potential of achieving a greater reduction in bit rate than the first method,  
9 because each escape sequence can now be replaced by the codeword for one special  
10 symbol, rather than by the two codewords as is the case for the first method. The second  
11 method, however, may reduce the PSNR due to increased quantization noise resulting  
12 from the process producing the lower levels. Therefore, if a desired reduction of escape  
13 sequences can be achieved using the first method, then there is no need to perform the  
14 second method, which is likely to reduce the PSNR. If the first method is used but not all  
15 of the escape sequences have been eliminated, then the second method could be used to  
16 possibly eliminate the remaining escape sequences.

17 The mapping of coefficient magnitudes to the levels can be changed by decoding  
18 the levels to coefficient magnitudes, changing the quantization scale factor (qsi), and then  
19 re-coding the levels in accordance with the new quantization scale factor (qsi). The  
20 quantization scale factor is initialized in each slice header and can also be updated in the  
21 macroblock header on a macroblock basis. Therefore it is a constant for all blocks in the  
22 same macroblock. In particular, the quantization scale factor is a function of a  
23 q\_scale\_type parameter and a quantizer\_scale\_code parameter. If q\_scale\_type = 0, then

1 the quantizer scale factor (qsi) is twice the value of q\_scale\_code. If q\_scale\_type =1,  
2 then the quantizer scale factor (qsi) is given by the following table, which is the right half  
3 of Table 7-6 on page 70 of ISO/IEC 13838-2:1996(E):

4

5	<u>quantizer scale code</u>	<u>quantization scale factor (qsi)</u>
6	1	1
7	2	2
8	3	3
9	4	4
10	5	5
11	6	6
12	7	7
13	8	8
14	9	10
15	10	12
16	11	14
17	12	16
18	13	18
19	14	20
20	15	22
21	16	24
22	17	28
23	18	32

1           19                           36  
2           20                           40  
3           21                           44  
4           22                           48  
5           23                           52  
6           24                           56  
7           25                           64  
8           26                           72  
9           27                           80  
10          28                           88  
11          29                           96  
12          30                           104  
13          31                           112

14  
15         In a preferred implementation, to reduce the coefficient levels, the quantization  
16         scale factor is increased by a factor of two, and the levels of the non-zero AC DCT  
17         coefficients are reduced by a factor of two, so long as the original value of the  
18         quantization scale factor is less than or equal to one-half of the maximum possible  
19         quantization scale factor. For q\_scale\_type = 1, a factor of two increase in the  
20         quantization scale factor (qsi) is most easily performed by a table lookup of a new  
21         quantization\_scale\_code using the following conversion table:

22

23         Original quantization\_scale\_code                           New quantization scale code

1	1	2
2	2	4
3	3	6
4	4	8
5	5	9
6	6	10
7	7	11
8	8	12
9	9	14
10	10	16
11	11	17
12	12	18
13	13	19
14	14	20
15	15	21
16	16	22
17	17	24
18	18	25
19	19	26
20	20	27
21	21	28
22	22	29
23	23	30

2

3 **V. Trick Mode Files**

4 In a preferred method for generation of trick mode files, the quantization scale  
5 factor is adjusted in order to achieve a desired reduction in the escape sequence  
6 occurrence frequency resulting from the modified FDSNR\_LM procedure, and the  
7 number (k) of largest magnitude coefficients is adjusted in order to achieve a desired  
8 reduction in bit rate. A specific implementation is shown in the flow chart of FIGS. 22-  
9 23. In a first step 361, the number (k) of largest magnitude AC coefficients per 8x8 block  
10 is initially set to a value of 9, and the quantization scaling factor (QSF) is initially set to a  
11 value of 2. Then conversion of the I frames of an original-quality MPEG-2 coded video  
12 clip to a lower quality level begins. When a picture header is encountered in step 362,  
13 indicating the beginning of a new I frame, execution continues to step 363. In step 363,  
14 execution branches depending on the value of the intra\_vlc\_format parameter in the  
15 picture header of the original-quality MPEG-2 coded video clip. This value is either 0,  
16 indicating that the first (run, level) coding table (TABLE 0) was used for coding the  
17 picture, or 1, indicating that the second (run, level) coding table (TABLE 1) was used for  
18 coding the picture. In either case, the down scaled quality picture will be coded with the  
19 second (run, level) coding table. If the intra\_vlc\_format parameter is equal to 0 execution  
20 continues from step 363 to step 364 where TABLE 0 is read in for (run, level) symbol  
21 decoding in the original-quality MPEG-2 coded clip. Otherwise, if the intra\_vlc\_format  
22 parameter is equal to 1, then execution continues from step 363 to step 365 where

2025 RELEASE UNDER E.O. 14176

1 TABLE 1 is read in for (run, level) symbol decoding in the original-quality MPEG-2  
2 coded clip.

3 After steps 364 and 365, execution continues to step 366. In step 366, the  
4 modified FDSNR\_LM procedure is applied to the 8x8 blocks of the current slice, using  
5 the adjusted quantization scale index, if the adjusted quantization scale index is less than  
6 the maximum possible quantization scale index. In step 367, execution loops back to step  
7 362 to continue 8x8 block conversion until a new slice header is encountered, indicating  
8 the beginning of a new slice. Once a new slice is encountered, execution continues from  
9 step 367 to step 368. In step 368, the average escape sequence occurrence frequency per  
10 block for the last slice is compared to a threshold TH1. If the escape sequence  
11 occurrence frequency is greater than the threshold, then execution branches to step 369.  
12 In step 369, if the quantization scaling factor (QSF) is less than or equal to a limit value  
13 such as 2, then execution branches to step 370 to increase the quantization scaling factor  
14 (QSF) by a factor of two.

15 In step 368, if the escape sequence occurrence frequency is not greater than the  
16 threshold TH1, then execution continues to step 371 of FIG. 23. In step 371, the average  
17 escape sequence occurrence frequency per 8x8 block for the last slice is compared to a  
18 threshold TH2. If the escape sequence occurrence frequency is less than the threshold  
19 TH2, then execution branches to step 372. In step 372, if the quantization scaling factor  
20 (QSF) is greater than or equal to a limit value such as 2, then execution branches to step  
21 373 to decrease the quantization scaling factor (QSF) by a factor of two. After step 373,  
22 and also after step 370 of FIG. 22, execution continues to step 374 of FIG. 23. In step  
23 374, execution continues to step 375 if a backtrack option has been selected. In step 375,

1 re-coding for the last slice is attempted using the adjusted quantization scale factor. The  
2 new coding, or the coding that gives the best results in terms of the desired reduction of  
3 escape sequence occurrence frequency, is selected for use in the scaled quality picture.  
4 After step 375, execution continues to step 376. Execution also continues to step 376  
5 from: step 369 in FIG 22 if the quantization scaling factor (QSF) is not less than or equal  
6 to 2; step 371 in FIG 23 if the escape sequence occurrence frequency is not less than the  
7 threshold TH2; step 372 in FIG 23 if the quantization scaling factor (QSF) is not greater  
8 than or equal to 2; and from step 374 in FIG 23 if the backtrack option has not been  
9 selected.

10 In step 376, the average bit rate of the (run, level) coding per 8x8 block for at  
11 least the last slice is compared to a high threshold TH3. Preferably this average bit rate is  
12 a running average over the already processed segment of the current scaled quality I-  
13 frame, and the high threshold TH3 is selected to prevent video buffer overflow in  
14 accordance with the MPEG-2 Video Buffer Verifier restrictions. If the average bit rate  
15 exceeds the high threshold TH3, then execution continues to step 377, where the number  
16 (k) of non-zero largest magnitude AC coefficients per 8x8 block is compared to a lower  
17 limit value such as 6. If the number (k) is greater than or equal to 6, then execution  
18 continues to step 378 to decrement the number (k).

19 In step 376, if the average bit rate is not greater than the threshold TH3, then  
20 execution continues to step 379. In step 379, the average bit rate is compared to a lower  
21 threshold TH4. If the average bit rate is less than the threshold TH4, then execution  
22 branches from step 379 to step 380, where the number (k) of non-zero largest magnitude  
23 AC DCT coefficients per 8x8 block is compared to a limit value of 13. If the number (k)

1       is less than or equal to 13, then execution continues to step 381 to increment the number  
2       (k). After step 378 or 381, execution continues to step 382. In step 382, execution  
3       continues to step 383 if a backtrack option is selected. In step 383, an attempt is made to  
4       re-code the last slice for the scaled quality picture using the adjusted value of the number  
5       (k) of non-zero largest magnitude AC DCT coefficients per block. After step 383,  
6       execution loops back to step 362 of FIG. 22 to continue generation of the scaled quality  
7       clip. Execution also loops back to step 362 of FIG. 22 after: step 377 if the value of (k) is  
8       not greater than or equal to 6; step 379 if the average bit rate is not less than the threshold  
9       TH4; step 380 if the value of (k) is not less than or equal to 13; and step 382 if the  
10      backtrack option has not been selected. Coding of the scaled quality clip continues until  
11      the end of the original quality clip is reached in step 384 of FIG. 22, in which case  
12      execution returns.

13           In a preferred implementation, a fast forward trick mode file and a fast reverse  
14      trick mode file are produced from an original-quality MPEG-2 coded video main file  
15      when the main file is ingested into the video file server. As shown in FIG. 24, a volume  
16      generally designated 390 is allocated to store the main file 391. The volume 390 includes  
17      an allocated amount of storage that exceeds the real file size of the main file 391 in order  
18      to provide additional storage for meta-data 392, the fast forward trick file 393, and the  
19      fast reverse trick file 394. The trick files are not directly accessible to clients as files;  
20      instead, the clients may access them through trick-mode video service functions. With  
21      this strategy, the impact on the asset management is a minimum. No modification is  
22      needed for delete or rename functions.

1       Because the volume allocation is done once for the main file and its fast forward  
2 and fast reverse trick mode files, there is no risk of lack of disk space for production of  
3 the trick files. The amount of disk blocks to allocate for these files is computed by the  
4 video service using a video service parameter (vsparams) specifying the percentage of  
5 size to allocate for trick files. A new encoding type is created in addition to types RAW  
6 for direct access and MPEG2 for access to the main file. The new encoding type is called  
7 EMPEG2, for extended MPEG2, for reference to the main file plus the trick files. The  
8 video service allocates the extra file size only for these files.

9           For the transfer of these files to archive or to another video file server, it would be  
10 useful to transfer all the data even if it is a non-standard format. For the FTP copy-in, a  
11 new option is added to specify if the source is in the EMPEG2 format or if it is a standard  
12 MPEG2 file. In the first case, the copy-in should provide the complete file 390. In the  
13 second case, the video service allocates the extra size and the processing is the same as  
14 for a record. For the copy-out, the same option can be used to export the complete file  
15 390 or only the main part 391. The archiving is always done on the complete file 390.

16           The trick mode file production is done by a new video service procedure. This  
17 procedure takes as input the speed-up factor (or the target trick mode file size) along with  
18 the number of freeze (P or B) frames to insert in between the scaled I frames and then  
19 generates both the fast forward file 393 and the fast reverse file 394 for this speed-up  
20 factor (or target trick mode file size) and with the specified number of interleaving freeze  
21 frames. Since the bandwidth of the original clip (in the main file) and the bandwidths of  
22 the two trick mode clips (in the fast forward and fast reverse files) are the same, the  
23 speed-up factor and the target trick mode file size are equivalent pieces of information. A

1 default speed-up factor (system parameter) can be used. The main file is read and the  
2 trick mode files are produced. If a trick mode file already exists with the same speed-up  
3 factor, it is rewritten or nothing is done depending on an option. Multiple trick mode  
4 files could be created with different speed-up factors. But it is preferred to permit only  
5 one set of fast forward and fast reverse trick mode files to be produced at a time (i.e., no  
6 parallel generation with different speed-up factors). The current speed-up factor is a  
7 parameter within the video service parameters (vsparams).

8 As stated above another parameter to be provided to the video service procedure  
9 in charge of trick mode file generation is the number of freeze frames to be inserted in  
10 between consequent scaled I frames. The preferred values for this parameter are 0 and 1,  
11 although other positive integer values greater than 1 are also possible. The inclusion of  
12 freeze frames due to their very small sizes spare some bandwidth which can then be used  
13 to improve the quality of scaled I frames. Hence, the freeze frames in this context provide  
14 a mechanism to achieve a trade-off between the scaled I frame quality and the temporal  
15 (motion) sampling. Depending on the speed-up factor (or the target trick mode file size)  
16 and also the number of interleaving freeze frames to be inserted, the video service  
17 procedure in charge of trick mode file generation determines a sub-sampling pattern  
18 (closest to uniform) to choose the original I frames which will be scaled and included in  
19 the trick mode files. For example, the case of an original clip with 10 frames per GOP, a  
20 trick mode file size which is 10% of the main file together with 0 freeze frames, implies  
21 the use of all original I frames for being scaled and included in the trick mode file. This  
22 will typically result in a low quality scaling. As another example, the case of an original  
23 clip with 10 frames per GOP, a trick mode file size which is 10% of the main file

1 together with 1 freeze frame, implies the use of a 2 to 1 (2:1) sub-sampling on the  
2 original I frames which will choose every other original I frame for being scaled and  
3 included in the trick mode file.

4 FIG. 25 is a more detailed diagram of the volume 390, showing additional meta-  
5 data and related data structures. The Inode 401 includes 4 disk blocks containing a file-  
6 system oriented description of the file. The meta-data (MD) directory 402 includes 4  
7 disk blocks describing each entry of the meta-data area 392. The entries of the meta-data  
8 area 392 include a description of the MPEG-2 meta-data 403, a description of the trick  
9 files header meta-data 404, and a description of the GOP index meta-data 405. The  
10 MPEG-2 meta-data 403 includes 15 disk blocks maximum.

11 The trick files header 404 includes 1 disk block, which specifies the beginning of  
12 free area (end of last trick file) in blocks, the number of trick files couple (FF FR), and  
13 for each trick file, a speed-up factor, a block address of the GOP index, a block address of  
14 the trick file forward, a byte length of the trick file forward, a block address of the trick  
15 file reverse, a byte length of the trick file reverse, a frames number of the trick file, and a  
16 number of GOP of each trick files.

17 The GOP index includes 2024 disk blocks. The GOP index specifies, for each  
18 GOP, a frame number, a pointer to the MPEG-2 data for the GOP in the main file, and  
19 various flags and other attributes of the GOP. The flags indicate whether the GOP entry  
20 is valid and whether the GOP is open or closed. The other attributes of the GOP include  
21 the maximum bit rate, the average bit rate, the AAU size in bytes, the APU duration in  
22 seconds, the audio PES packet starting locations, the AAU starting locations, the AAU  
23 PTS values, and the decode time stamp (DTS) and the value of the program clock

1 reference (PCR) extrapolated to the first frame of the GOP. The size of all the data  
2 preceding the main file is, for example, 1 megabyte.

3 There is one GOP index 406 for both the fast forward file 393 and the fast reverse  
4 file 394. The GOP index 406 of the trick files is different than the GOP index 405 of the  
5 main file. The GOP index 406 of the trick files contains, for each GOP, the byte offset in  
6 the trick file forward of the TS packet containing the first byte of the SEQ header, the  
7 frame number in the fast forward file of the GOP (the same value for the fast reverse file  
8 can be computed from this value for the fast forward file), the frame number in the  
9 original file of the first frame of the GOP, and the byte offset in the original file of the  
10 same frame (to resume after fast forward or reverse without reading the main GOP  
11 index).

12 The GOP index 405 for the main file and the GOP index 406 for the fast forward  
13 and fast reverse trick files provide a means for rapidly switching between the normal  
14 video-on-demand play operation during the reading of the main file, and the fast-forward  
15 play during the reading of the fast-forward file, and the fast-reverse play during the  
16 reading of the fast-reverse file. For example, FIG. 26A illustrates the read access to  
17 various GOPs in the main file, fast forward file, and fast reverse file, during a play  
18 sequence listed in FIG. 26B. Due to the presence of down-scaled I frames and possibly  
19 present consequent freeze frames in the trick mode files, the video buffer verifier (VBV)  
20 model for a trick mode file is different than the VBV model of the main file.  
21 Consequently, the mean video decoder main buffer fullness levels can be significantly  
22 different for these files. For example, a transition from the main file to one of the trick  
23 files will usually involve a discontinuity in the mean video decoder main buffer fullness

1 level, because only the I frames of the main file correspond to frames in the trick files,  
2 and the corresponding I frames have different bit rates when the trick mode I frames are  
3 scaled down for a reduced bit rate. An instantaneous transition from a trick file back to  
4 the main file may also involve a discontinuity especially when freeze frames are inserted  
5 between the I frames for trick mode operation. To avoid these discontinuities, the  
6 seamless splicing procedure of FIGS. 3 to 6 as described above is used during the  
7 transitions from regular play mode into trick mode and similarly from trick mode back  
8 into the regular play mode. Through the use of the seamless splicing procedure to modify  
9 the video stream content, for example for the “Seamless Splice” locations identified in  
10 FIG. 26A, the video decoder main buffer level will be managed so as to avoid both  
11 overflows and underflows leading to visual artifacts.

12 It is desired to copy in and out of the volume 390 with or without the meta-data  
13 392 and the trick files 393, 394. This is useful to export and/or import complete files  
14 without regenerating the trick files. The file encoding type is now recognized as a part of  
15 the volume name. Therefore there can be multiple kinds of access to these files. The  
16 read and write operations are done by derivations of the class file system input/output  
17 (FSIO) which takes into account the proper block offset of the data to read or write.  
18 There is one derivation of FSIO per encoding type, providing three different access  
19 modes. EMPEG2, MPEG2, and RAW. EMPEG2 accesses the whole volume from the  
20 beginning of the meta-data array, and in fact provides access to the entire volume except  
21 the inode 401, but no processing is done. MPEG2 access only the main part of the asset  
22 with MPEG processing, including file analysis and meta-data generation in a write  
23 access. RAW access only the main part of the asset without processing. These access

1 modes are operative for read and write operations for various access functions as further  
2 shown in FIG. 27.

3 During a record operation, the video service allocates a volume and computes the  
4 number of blocks to allocate using the volume parameter giving the percentage to add for  
5 the trick files. Then, the size in blocks given to the stream server is the main part size  
6 only without the extension for the trick files. This avoids using the reserved part of the  
7 volume when the effective bit rate is higher than the requested bit rate. At the end of a  
8 record operation or an FTP copy-in operation, the video service calls a procedure  
9 CMSPROC\_GETATTR, and the stream server returns the actual number of bytes  
10 received and the actual number of blocks used by the main file plus the meta-data. The  
11 same values are returned for both MPEG2 and EMPEG2 files. The video service  
12 computes again the file extension to manage the trick files and adjust the number of  
13 allocated blocks.

14 Both trick files, forward and reverse, are generated by the same command. First,  
15 the trick file forward is generated by reading the main file. The trick file GOP index is  
16 concurrently built and kept in memory. During this generation, only the video packets  
17 are kept. PCR, PAT and PMT will be regenerated by the MUX in play as for any other  
18 streams. The audio packets are discarded. This ensures that there is enough stuffing  
19 packets for the PCR reinsertion. For this, a stuffing packet is inserted every 30  
20 milliseconds.

21 Then using the GOP index, the trick file forward is read GOP by GOP in reverse  
22 order to generate the trick file reverse. The same GOPs are present in both files. The  
23 only modification done is an update of the video PTS, which must be continuous. Then,

1 the GOP index is written on disk. This avoids reading again the file while generating the  
2 second trick file. The GOP index size is: 24 times the GOP number. In the worst case  
3 (the file is assumed not to be 1 frame only), there are 2 frames per GOP and 30 frames  
4 per second. So for 1 hour in fast forward, the GOP index size is:  $(24 \times 3600 \times 30) / 2 =$   
5 1296000 bytes. This will be the case for a 4 hour film played at 4 times the normal  
6 speed. Therefore, this GOP index can be kept in memory during the trick file generations  
7 without risk of memory overflow.

8       The read and write rates are controlled to conserve bandwidth on the cached disk  
9 array. The bandwidth reserved for these operations is a parameter given by the video  
10 service. It is a global bandwidth for both read and writes. The number of disk I/Os per  
11 second is counted so as not to exceed this bandwidth.

12       The trick files' header update is done once when both the fast forward and fast  
13 reverse trick files and the GOP index have been successfully written.

14       Playing a file is done with the CM\_MpegPlayStream class. Fast forward  
15 (reverse) can only be requested when the stream is in the paused state. The current frame  
16 on which the stream is paused is known from the MpegPause class. This frame is located  
17 in the GOP index of the trick file. Then the clip start point and length are modified in the  
18 Clip instance with the trick file position computed from the beginning of the clip. So, the  
19 Clip class handles these trick files in a manner similar to the main file. The current  
20 logical block number is updated with the block address in the trick file recomputed from  
21 the beginning of the main clip. In fact, a seek is performed in the trick file as it was part  
22 of the main file, which is totally transparent for the ClipList and Clip classes. The  
23 transition from fast forward to pause is handled in a similar fashion. The clip start and

1 length and the logical block number are again updated. The smooth transitions from  
2 pause to fast forward and from fast forward to pause are done in the same way as for  
3 regular play. There is a splicing from the pause stream to the play stream.

4 The class hierarchy for trick file handling is shown in FIG. 28. The MpegFast,  
5 MpegFastForward and MpegFastReverse class handles the GOP generation from the  
6 initial file. This is the common procedure for building the GOP whatever the source and  
7 the destination are. RealTimeFastFwd and RealTimeFastRev are the classes instantiated  
8 when a real time fast forward (reverse) has to be done. They manage the real-time buffer  
9 flow to the player. There is a derivation of the methods takeBuffer and returnBuffer  
10 which use the base class to build the GOP in the buffer to be played. The main file  
11 access is done using a buffer pool.

12 TrickFilesGenerate is the class instantiated to generate trick files forward and  
13 reverse. It inherits from TrickFilesAccess the methods for reading the original file into  
14 some buffers and for writing the trick file and its meta-data. It inherits from  
15 MpegFastForward the methods for building the GOP and for managing the advance in  
16 the file.

17 The computation of the next 1 frame to play is done by MpegFast,  
18 MpegFastForward and RealTimeFastFwd. When a trick file generation command is  
19 invoked, a thread is created and started and the generation itself is done off-line. A call-  
20 back is sent to the video service when the generation is completed. The class  
21 TrickFilesGenerate generates the trick file forward, and then, using the GOP index built  
22 in memory, the class TrickFilesGenerate generates the trick file reverse.

1        When there is a transition from play to pause, the only latency issue is related to  
2        the buffer queue handled by the player and to the GOP size. The stream can build  
3        immediately the active pause GOP, and then this GOP will be sent at the end of the  
4        current GOP with a splicing between these two streams.

5        When there are transitions from pause to regular play or fast forward and fast  
6        reverse, a seek in the file is done. This means that the current buffer pool content is  
7        invalidated and the buffer pool is filled again. Play can start again while the buffer pool  
8        is not completely full, as soon as the first buffer is read. The buffer pool prefilling can  
9        continue as a background process. The issue here is that there is a risk to generate an  
10      extra load on the cached disk array as well as on the stream server side when the buffer  
11      pool is being prefilled.

12       To avoid too frequent transitions from play to fast forward and fast reverse, there  
13      is a limitation of the number of requests per second for each stream. This limitation is  
14      part of the management of the video access commands. A minimum delay between two  
15      commands is defined as a parameter. If the delay between a request and the previous one  
16      is too small, the request is delayed. If a new request is received during this delay, the  
17      new request replaces the waiting one. So the last received request is always executed.

18       The video service parameters file (vsparams) contains these new parameters for  
19      the trick mode files:

20       TrickFileExtensionSize:<percent>:

21       DefaultFastAcceleration:<acceleration>:

1 DMtrickFileGen:<mask of reserved DM> (This parameter is a mask of the stream  
2 servers that can be chosen to perform the trick file generation. The default value is  
3 0xffffc: all of the stream servers.)

4 DMtrickFileGenBW:<bandwidth used for trick file generation> (This parameter  
5 is the value of the bandwidth effectively used by the stream server for the trick files  
6 generation.)

7 The video service routines are modified to operate upon the EMPEG2 files, and in  
8 particular to compute the size of the EMPEG2 files, to allocate the volume for the main  
9 file and the trick files, and to generate the trick files. The volume creation functions  
10 (VAPP) and volume access functions (RRP) use the EMPEG2 files in the same way as  
11 MPEG2 files. This means that an MPEG2 volume is created on the stream server. Both  
12 MPEG2 and EMPEG2 files can be used in the same session or play-list. The session  
13 encoding type is MPEG2. In record (or copy-in), the number of blocks allocated for an  
14 EMPEG2 file is computed using the percentage of size to add. At the end of record (or  
15 copy-in), the number of blocks is adjusted using the number of blocks returned by the  
16 stream server (by CMSPROC\_GETATTR) and adding the percentage for trick files. The  
17 trick files validity and generation date are stored by the video service in the asset  
18 structure. The bandwidth allocated to the TrickFilesGenerate command is defined in the  
19 video service parameters (vsparams or vssiteparams). The selection of a stream server to  
20 generate the trick files takes into account this bandwidth only. If preferred stream servers  
21 are specified in vsparams (or vssiteparams), then the selected stream server will be one of  
22 these specified stream servers.

1        In a preferred implementation of the video service software, a new encoding type  
2        is created. The encoding type enum becomes:

```
3        enum encoding-t{  
4            ENC_UNKNOWN = 0,            /* unknown format */  
5            ENC_RAW         = 1,            /* uninterpreted data */  
6            ENC_MPEG1      = 2,            /* constrained MPEG1 */  
7            EMC_MPEG      = 3,            /* generic MPEG */  
8            ENC_EMPEG2    = 4,            /* MPEG2 with trick files extension */  
9       };
```

10  
11 -        The encoding information accessible by VCMP\_EXTENDEDINFO includes  
12        information about trick files:

```
13  
14        struct trickFilesInfo_t{  
15            ulong_t        generationDate;        /* date/time of the generation of the trick  
16                          files */  
17            rate_factor_t acceleration;        /* acceleration factor */  
18            ulong_t        framesNumber;        /* frames number in each trick file (FWD and  
19                          REV) */  
20            ulong_t        gopNumber;        /* GOP number of each file */  
21       };  
22  
23        struct EMPEG2info_t{
```

```
1      MPEG2info_t      MPEG2info;
2      trickFileInfo_t   trickFiles<>;
3  };
4
5  union encodingInfo_t switch (encoding-t enc){
6      case ENC_MPEG:
7          MPEG2info_t      MPEG2info;
8      case ENC_EMPEG2:
9          EMPEG2info_t     EMPEG2info;
10     default:
11         void;
12     };
13
14 The video service software includes a new procedure (VCMP_TRICKFILESGEN) for
15 trick file generation, which uses the following structures:
16
17 struct VCMPtrickgenres_t{
18     VCMPstatus_t      status;
19     tHandle_t         handle;
20 };
21
22 struct VCMPtrickfilesargs_t{
23     name_t            clipname;
```

```
1     bool_t      overwriteIfExists;
2
3     rate_factor_t acceleration;
4
5     };
6
7
8     VCMPtrickgenres_t      VCMP_TRICKFILESGEN (VCMPtrickfilesargs_t)  =
9     36,
```

7

8        If the trick files already exist and if the boolean overwriteIfExists is true, then the  
9        trick files are generated again, in the other case nothing is done. Acceleration is the  
10      acceleration as defined and used for the controlled speed play function. It is a percentage  
11      of the normal speed, it must be greater than 200 and smaller than 2000. The special value  
12      of 0 can be used to generate files with the default acceleration defined in vssiteparams.  
13      The procedure starts the generation process. The completion is notified by a callback.

14      The video service includes a new option to copy-in and copy-out. The option is  
15      added to allow a user to copy all the file or the main asset only. For compatibility with  
16      old client programs, the following new procedures are added:

17

```
18     VCMPcopyres_t      VCMP_FULL_COPYIN    (copyinargs2_t)      = 37,
19     VCMPcopyres_t      VCMP_FULL_COPYOUT   (copyoutargs2_t)     = 38,
```

20

21      These new procedures have the same interface as the already existing one, but are used to  
22      copy-in the complete file: meta-data + asset + trick files.

23

1           The video service includes a new procedure

2       VCMP\_TRICKFILESGENCOMPLETED, which uses the following structures:

3

4       struct VCMPtrickfilecomplete\_t{

5           tHandle\_t                  handle;

6           VCMPstatus\_t              status;

7       };

8

9       VCMPstatus\_t TRICKFILESGENCOMPLETED (VCMPtrickfilecomplete\_t) = 10,

10

11          The video service includes new procedures added for handling trick mode

12       generation arguments and using the following structures:

13

14       struct cms\_trick\_gen\_args {

15           Handle\_t                 Vshandle;

16           name\_t                 name;

17           bool\_t                 overwritelfExists;

18           rate\_factor\_t         acceleration;

19           bandwidth\_t            reservedBw;

20       };

21

22       cms\_status CMSPROC\_GEN\_TRICK\_FILES (cms\_trick\_gen\_args) = 34,

23

```
1 struct trick_gen_completed_args {  
2     Handle_t          Vshandle;  
3     cms_status        status;  
4 };  
5 void CTLPROC_TRICKGENCOMPLETED (trick_gen_completed_args)      = 8,  
6
```

7       The video service includes the following option to force the regeneration of trick  
8       files even if they exist:

```
9           nms_content -gentrick <name> [<-f>] [acceleration]
```

10      Without this option, an error code is returned if the trick files exist. “Acceleration” is an  
11      acceleration factor. If it is not present, the default value is taken from vsparams.

12      The video services include an encoding information access function (nms\_content  
13      -m). This function produces displayed output containing, for each trick file generated,  
14      the acceleration, the generation date and time, the number of frames, and the number of  
15      GOPs.

16      For the use of an FTP copy function with the trick files, the following new  
17      commands are added:

18

```
19      nms_content -copyinfull <same arguments as -copyin>
```

```
20      nms_content -copyoutfull <same arguments as -copyout>
```

21

22      VI. Reduction of MPEG-2 Transport Stream Bit Rate for Combining Multiple  
23      MPEG-2 Transport Streams

1 Another application of the SNR scaling achieved by the invention is to reduce the  
2 bit rate of an MPEG-2 transport stream in order to allow combining multiple MPEG-2  
3 transport streams to match a target bit rate for a multiple program transport stream. For  
4 example, FIG. 29 shows a system for combining an MPEG-2 audio-visual transport  
5 stream 411 with an MPEG-2 closed-captioning transport stream 412 to produce a  
6 multiplexed MPEG-2 transport stream 413. In this case, the closed captioning transport  
7 stream 412, containing alphanumeric characters and some control data instead of audio-  
8 visual information, has a very low bit rate compared to the audio-visual transport stream  
9 411. Assuming that the target bit rate for the multiplexed transport stream 413 is the  
10 same as the bit rate of the audio-visual transport stream 411, there need be only a slight  
11 decrease in the bit rate of the audio-visual transport stream, and this slight decrease can  
12 be obtained by occasionally removing one non-zero AC DCT coefficient per 8x8 block.  
13 Therefore, in the system of FIG. 29, the audio-visual transport stream 411 is processed by  
14 a program module 414 for selective elimination of non-zero AC DCT coefficients to  
15 slightly reduce the average bit rate of this transport stream. A transport stream  
16 multiplexer 415 then combines the modified audio-visual transport stream with the closed  
17 captioning transport stream 412 to produce the multiplexed MPEG-2 transport stream  
18 413.

19 In order to determine whether or not any non-zero AC DCT coefficient should be  
20 eliminated from a next 8x8 block in the audio-visual transport stream 411, a module 421  
21 is executed periodically to compute a desired bit rate change in the audio-visual transport  
22 stream 411. For example, respective bit rate monitors 416, 417 may measure the actual  
23 bit rate of the audio-visual transport stream 411 and the closed captioning transport

1 stream 412. Alternatively, if it is known precisely how these transport streams are  
2 generated, presumed values for the bit rates of these transport streams may be used in lieu  
3 of measured bit rates. The computation of the desired bit rate change also includes the  
4 desired bit rate 418 for the multiplexed MPEG-2 transport stream, and a bit rate 419 of  
5 multiplexer overhead, representing any net increase in bit rate related to the multiplexing  
6 of the audio-visual transport stream 411 with the closed captioning transport stream 412.  
7 An adder/subtractor 420 combines the various bit rate values from the inputs 416, 417,  
8 418, and 419 to compute the desired bit rate change in the audio-visual transport stream  
9 411. From the adder/subtractor 420 output, the module 421 converts the desired change  
10 in bit rate to a desired number of bits to be removed per computational cycle (e.g., per  
11 millisecond). This number of bits to be removed per computational cycle is received in  
12 an adder/subtractor 422, and the output of the adder/subtractor is received in an integrator  
13 423. A limiter 424 takes the sign (positive or negative) of the integrated value to produce  
14 a flag indicating whether or not one non-zero AC DCT coefficient should be removed  
15 from the coefficients for the next 8x8 block, assuming that the next block has at least one  
16 non -zero AC DCT coefficient. (Alternatively, a non-zero AC DCT coefficient could be  
17 removed only if the 8x8 block has more than a predetermined fraction of the average  
18 number of AC DCT coefficients per 8x8 block.) The particular non-zero AC DCT  
19 coefficient to remove in each case can be selected using any of the methods discussed  
20 above with reference to FIGS. 14, 15, or FIG. 20. For example, the coefficient to remove  
21 could be the last non-zero AC DCT coefficient in the scan order. Alternatively, the non-  
22 zero AC DCT coefficient having the smallest magnitude could be removed so long as its  
23 removal does not cause an escape sequence.

1 When the module 414 removes a non-zero AC DCT coefficient from a 8x8 block,  
2 it sends the number of bits removed to the adder/subtractor 422. In a preferred  
3 implementation, the operations of the adder/subtractor 422, integrator 423, and limiter  
4 424 are performed by a subroutine having a variable representing the integrated value.  
5 During each computational cycle, the variable is incremented by the number of bits to be  
6 removed per computational interval, and whenever the module 414 removes a non-zero  
7 AC DCT coefficient from a 8x8 block of the audio-visual transport stream, the variable is  
8 decremented by the number of bits removed.

9 Although the system in FIG. 29 has been described for achieving a slight  
10 reduction in bit rate of the MPEG-2 audio-visual transport stream 411 for combining  
11 multiple transport streams to produce a multiple program MPEG-2 transport stream, it  
12 should be apparent that it could be used for obtaining relatively large reductions in bit  
13 rate. In this case, the module 414 would use the procedure of FIGS. 14, 15 or preferably  
14 FIG. 20, and a multi-level comparator 424 would be used instead of a single-level  
15 comparator 424. The multi-level comparator would determine a desired number of non-  
16 zero coefficients to discard per 8x8 block based on the value of the output of the  
17 integrator 423. The maximum number of non-zero AC coefficients to keep for each 8x8  
18 block (i.e., the value of the parameter "k"), for example, would be determined by  
19 subtracting the number of non-zero AC DCT coefficients in the 8x8 block from the  
20 desired number to discard, and limiting this difference to no less than a predetermined  
21 fraction of the average number of non-zero AC coefficients per 8x8 block.

22

1           VII. Largest Magnitude Indices Selection For (Run, Level) Encoding Of A Block

2       Coded Picture

3           As described above with reference to FIG. 15, one way of scaling original-quality  
4           MPEG-2 video to produce lower-quality MPEG video is to retain up to a certain number  
5           of largest-magnitude non-zero AC DCT coefficients and to truncate any remaining non-  
6           zero AC DCT coefficients from each 8x8 block of pixels. This was referred to as the  
7           FDSNR\_LM procedure. As shown and described with reference to FIG. 15, in step 262,  
8           the (run, level) event variable-length codes (VLCs) for each block are parsed and  
9           decoded to produce a set of quantization indices. In step 263, the quantization indices  
10          (for AC DCT coefficients) are transformed to quantized coefficient values, and in step  
11          264, the quantized coefficient values are sorted in descending order of their magnitudes.

12          In order to perform scaling in a computationally more efficient way, it is possible  
13          to eliminate the step 263 of transforming the quantization indices to quantized coefficient  
14          values by selecting largest magnitude quantization indices (for the AC DCT coefficients)  
15          instead of selecting the largest magnitude quantized coefficient values. This would make  
16          the scaling procedure more suitable for real-time applications, so long as there would not  
17          be a significant degradation in performance compared to the performance obtained by  
18          selecting largest magnitude quantized coefficient values. For comparison purposes, the  
19          method of selecting the largest magnitude quantization indices will be referred to as  
20          “LMIS” (Largest Magnitude Indices Selection) and the method of selecting the largest  
21          magnitude quantized coefficient values will be referred to as “LMCS” (Largest  
22          Magnitude Coefficient Selection). It has been discovered that LMIS is not only

1 computationally more efficient than LMCS but also LMIS provides an improvement in  
2 performance over LMCS in the rate-distortion sense.

3 FIG. 30 shows the LMIS procedure performed on an 8x8 pixel block for the case  
4 where no pivoting is used. (The use of pivoting will be discussed below with reference to  
5 FIG. 43.) In other words, a subroutine corresponding to the flowchart in FIG. 30 is called  
6 once for each series of variable-length codes corresponding to an 8x8 pixel block. In a  
7 first step 461, the differential DC coefficient representing variable-length code (VLC) is  
8 parsed and copied from an input bit stream of original-quality MPEG-2 video to an  
9 output bit stream of reduced-quality MPEG video. Then in step 462, all of the following  
10 (run, level) event variable-length codes (VLCs) are parsed and decoded until and  
11 including the first end-of-block (EOB) marker in the input bit stream. In step 463, the  
12 quantization indices are sorted in descending order of their magnitudes. Step 463 could  
13 use any of the sorting methods described above with reference to FIGs. 14 to 19. In step  
14 464, up to the first K indices of the sorted list are kept and the last 63-K indices of the  
15 sorted list are in effect set to zero. In other words, the last 63-K indices in the sorted list  
16 are not allowed to be represented by nonzero levels of (run, level) events in the output bit  
17 stream for the lower-quality MPEG video, but rather these last 63-K indices in the sorted  
18 list contribute to runs of zeros. In step 465, (run, level) event formation and entropy  
19 encoding is applied to the new set of up to the first K indices in the sorted list. In the last  
20 step 466, the resulting VLCs are copied to the output bit stream until and including the  
21 end of block (EOB) marker.

22 The performance of LMIS would be equivalent to the performance of LMCS if  
23 the quantization of the coefficient values affected all coefficients to the same degree.

1 However, MPEG-2 provides a visually weighted quantization matrix that modifies the  
2 quantization step size within a block. (See, for example, Section 7.4, Inverse  
3 Quantization, page 68, of the MPEG-2 International Standard ISO/IEC 13818-2  
4 Information technology - Generic coding of moving pictures and associated audio  
5 information: Video, 1996.) Therefore, DCT coefficients at higher frequencies with larger  
6 magnitudes may be mapped to indices with smaller values than DCT coefficients at lower  
7 frequencies with smaller magnitudes. Consequently, the ordering of the largest  
8 magnitude quantized coefficient values (produced by sorting the coefficient magnitudes  
9 in step 264 of FIG. 15) is not necessarily the same as the ordering of the largest  
10 magnitude quantization indices (produced by sorting index magnitudes in step 463 of  
11 FIG. 30).

12 If the video tends to have a predominant high spatial frequency content, then the  
13 coefficient matrix will tend to have larger values for the higher frequencies than the lower  
14 frequencies, and the visually weighted quantization matrix will cause the LMIS  
15 procedure to favor the selection of lower-frequency coefficients. In this case, the LMIS  
16 procedure will trade many of the indices with value one in the mid-to-high frequency  
17 range with value one (and sometimes 2 or 3) indices in the low frequency range. This  
18 swap of indices may result in lower picture signal-to-noise ratio (PSNR) than that  
19 obtained via the LMCS procedure for the same number of retained AC coefficients.  
20 However, the bit savings achieved by the LMIS procedure are much more significant  
21 than the PSNR loss, and overall, the performance of LMIS is a lot better in the rate-  
22 distortion sense than that of LMCS. The indices that are retained by the LMIS procedure  
23 will not only shorten the run-lengths but also from the perceptual coding point of view, it

1 will result in more pleasant images better matching subjectively/visually the Human  
2 Visual System's (HVS) low-pass resembling nature. That is not to say that mid-to-high  
3 frequency components are not important. It means simply that when the coefficient  
4 values are comparable to a certain extent, it is a better decision to keep the ones in the  
5 low-pass band both for bit savings and for achieving perceptual coding. Even though  
6 when two indices at different frequency channels have the same magnitude, the  
7 corresponding coefficient at the higher frequency channel being in general larger than  
8 that at the lower frequency channel, the difference does not justify the significantly  
9 higher bit-rate cost of coding the larger coefficient.

10 It may also be argued that LMIS is almost the same as the low-pass (LP) scaling  
11 described above with reference to FIG. 14. This is true for blocks of low-pass nature, and  
12 all three procedures (LMCS, LMIS, and LP) behave more or less the same when the  
13 signal power is concentrated in the low-pass band. However, when there is high signal  
14 power present in some other frequency band(s) as opposed to or in addition to the low-  
15 pass band, it is usually desirable to keep the signal components in those bands. While the  
16 LP scaling procedure can not achieve this, the LMCS procedure, being sensitive even to  
17 the smallest difference in the coefficient values, retains such bands even when the power  
18 in those bands is comparable to the power in the lower-frequency bands. Not only due to  
19 high coding cost of indices in such bands but also due to the low-pass nature of the  
20 human visual system, we would like to favor the lower-frequency bands to higher  
21 frequency bands when the power levels are comparable. In many cases where the LMCS  
22 procedure favors the indices in higher frequency bands, it may be better to favor the  
23 indices in low-frequency bands due to extremely higher cost of coding indices in higher

frequency bands. Two intermingled effects are present here. First, the indices at the higher frequencies are usually paired with longer run lengths, and secondly, the exclusion of indices in lower frequencies will result in even longer run lengths for the indices to be kept in those bands. Therefore, in the rate-distortion sense, the signal components in higher frequencies should be retained when they indeed represent some significantly strong image feature. LMIS does that very well. It balances the power and bit budget and hence results in better rate-distortion curves. LMCS, on the other hand, is strictly focused to obtaining the most signal power without any attention to the bit budget. If the bid budget were not a concern, LMCS could be a better choice. However, the problem domain dictates that these two factors must be balanced in an appropriate manner.

FIGS. 31 and 32 show performance comparisons between the LMCS and LMIS procedures for the case of no pivot insertion and the use of MPEG-2's default visually weighted quantization matrix having larger values for the higher frequencies than the lower frequencies for video. The performance improvement of LMIS over LMCS is less significant when pivot insertion is used, as described below, for avoiding escape sequences and reducing the bits needed for (run, level) coding.

17

### 18        VIII. Avoiding Escape Sequences In Coding of (Run, Level) Pairs

19        In view of the above, there have been described methods of efficient SNR scaling  
20        of video originally present in a high-quality and nonscalable MPEG-2 transport stream.  
21        To reduce the bandwidth of nonscalable MPEG-2 coded video, certain non-zero AC DCT  
22        coefficients for the 8x8 blocks are removed from the MPEG-2 coded video.

1 It is recognized that the largest magnitude coefficient selection (LMCS) procedure  
2 of FIG. 15, in theory, has a great potential to provide high-quality scaling. However, in  
3 practice, under the practical bit rate versus quality, peak signal-to-noise (PSNR) rate-  
4 distortion measure, the LMCS procedure has a performance problem. The source of the  
5 problem, in the most part, appears to be a mismatch (i.e. a non-compatibility), between  
6 the (run, level) event statistics generated by the LMCS procedure and the statistics that  
7 the MPEG-2 (run, level) VLC codebooks are designed for. This mismatch revealed itself  
8 by both the generation of a drastically increased number of escape sequences and also an  
9 increased tendency towards using less likely (according to the MPEG-2 base statistics)  
10 (run, level) symbols represented by longer code words.

11 There are two principal and coupled mechanisms leading to the problem. First of  
12 all, the LMCS procedure, by its very nature, retains the larger magnitude coefficients.  
13 Secondly, smaller magnitude coefficients in between the larger ones (to be retained) are  
14 discarded, leading to longer run lengths. The MPEG VLC codebooks are such that the  
15 larger the magnitude of indices and the longer the run lengths, the more bits are required  
16 to code them. Very often, the (run, level) pairs generated via LMCS fall out of the  
17 codebook, necessitating to resort to the costly (fixed 24 bits) Escape Sequence coding,  
18 which is a mechanism to code very rarely, in a statistical sense, occurring (run, level)  
19 pairs.

20 As described above with reference to FIGS. 20 and 21, one way of avoiding  
21 escape sequences from the LMCS procedure was to include a non-zero, non-qualifying  
22 AC DCT coefficient in the (run, level) coding. As described below, this method can be  
23 extended to avoid escape sequences or reduce the number of bits used in the variable-

length coding by introducing indices of magnitude 1 into coefficient channels corresponding to zero-valued AC DCT coefficients in the original-quality MPEG-2 coded video. In any case, a quantization index introduced into the coding for the reduced-quality MPEG coded video for the purpose of avoiding an escape sequence or reducing the number of bits for variable-length coding will be referred to as a “pivot index.” The pivot indices, when used jointly with the LMCS or LMIS procedures, effectively change the original statistics of the (run, level) symbols generated by the baseline LMCS or LMIS procedures. In fact, the pivot technique, as described below, is useful in combination with any encoding or scaling technique producing (run, level) codes that deviate from the normal MPEG statistics by having a greater than normal frequency of escape sequences and an increased probability mass on (run, level) symbols with relatively long codewords.

The objective of inserting pivot indices is to break the long run-lengths of zero valued AC DCT coefficients when such a split leads to a savings in the number of bits required for encoding. The basic underlying principle is that if preserving a non-qualifying smaller magnitude coefficient normally to be discarded by the LMCS (or the LMIS) procedure requires fewer bits to encode both itself and the following larger magnitude coefficient which was originally to be retained, then one can shoot two birds with one stone by preserving this non-qualifying coefficient. That is to say, not only a bit savings is achieved but also the quality, i.e. the PSNR measure, improves due to the inclusion of one more genuine coefficient.

The following example illustrates the basic underlying principle. Assume that in a sample quantized 8x8 coefficient block, a partial listing of the indices ordered

1 according to the employed zigzag scan order is given as (... , 9, 0, 3, 6, ...). Assume  
2 further that the LMCS algorithm decides to retain the coefficients associated with the  
3 indices 9 and 6 but not 3. Then, the index 3 will be treated as zero resulting in a (run,  
4 level) pair of value (2, 6). The symbol (2, 6) is not allocated a particular variable length  
5 codeword in the codebook and hence its encoding requires the use of an Escape Sequence  
6 of 24 bits. However, if we decide to retain the index 3 as well, then two alternate (run,  
7 level) pairs, namely (1, 3) and (0, 6), need to be encoded instead. Since the encoding of  
8 the latter two symbols requires 15 bits in total, 9 bits are saved with respect to the first  
9 alternative. Also, the inclusion of the index 3 contributes in a positive sense to reduce the  
10 power of the reconstruction error. Here, the index 3 is the pivot index. This technique is  
11 the first version, called Pivot 1, of a class of pivot techniques summarized in FIG. 33. As  
12 shown in the first step 481 of FIG. 33, the Pivot-1 technique selectively retains genuine  
13 non-qualifying non-zero AC coefficients in order to avoid escape sequences.

14 The primary motivation for preserving the index 3 in the above example is to save  
15 bits by avoiding the generation of an escape sequence which is the most inefficient way  
16 of encoding quantized coefficient data in MPEG-2. The marginal improvement in the  
17 PSNR came as a side benefit. An analysis of both of the (run, level) VLC codebooks  
18 (Table 0 and Table 1) employed by MPEG-2 reveals the fact that for a fixed value of the  
19 run-length, the lengths of the codewords are always defined by a monotonic non-  
20 decreasing function of the level. Since the marginal SNR improvement provided by the  
21 pivot index is of secondary significance, why not, then, change the value of the pivot  
22 index to plus one or minus one depending on the sign of its original value and achieve a  
23 further savings in bits? Going back to our previous example, when we apply this idea to

1 the original two (run, level) pairs, we get the symbols (1, 1) and (0, 6) the encoding of  
2 both of which requires 11 bits instead of 15. It should be also noted that, even though to  
3 a lesser extent with respect to its preservation in its original value, the inclusion of an  
4 index with magnitude one and the correct sign, still contributes positively to the quality  
5 (i.e., PSNR) as compared to the case of its total elimination. This version will be called  
6 the Pivot-2 technique. As shown in step 482 of FIG. 33, the Pivot-2 technique reduces  
7 the level magnitudes of the retained non-qualifying non-zero AC coefficients to a value  
8 of one in order to eliminate more escape sequences and to reduce the number of bits for  
9 (run, level) encodings.

10 A third and final version of the pivot techniques, called Pivot-3, involves inserting  
11 a pivot of level magnitude 1 for a level zero coefficient in the transformation of the  
12 original high-quality 8x8 pixel block into the lower quality version. (See step 483 in  
13 FIG. 33.) In effect, the pivot is noise that is inserted into the picture to obtain a more than  
14 compensating benefit of reducing the number of bits to encode the picture. Moreover, the  
15 objective reduction in PSNR due to inserting the noise-like pivots is masked subjectively  
16 by inter-coefficient contrast masking and in many cases is not visible to a casual human  
17 observer.

18 Consider first the relation of the escape sequence count with the number of  
19 preserved coefficients in each block. FIGs. 34 and 35 show the plots of the average  
20 number of escape sequences per frame as a function of the number of LMCS coefficients  
21 retained in each block for various quantization levels (qsv) at the input of the LMCS  
22 processing, and with or without the various pivoting mechanisms. The data for these  
23 plots was produced by averaging over representative frames from the three standard test

1 sequences (namely, Susie, Flower Garden and Football). In these plots and for a fixed  
2 input quantization level, the general trend which is actually to a significant extent  
3 common to all four (including Pivot-3) illustrated different LMCS and pivoting  
4 combinations, is as follows. When the baseline LMCS algorithm is configured to  
5 preserve only a few largest magnitude coefficients, the number of escape sequences  
6 generated is quite high. This is not only because the preserved indices are (most of the  
7 time) the largest magnitude indices of all but also because of the fact that the few number  
8 of preserved indices achieve only a very sparse sampling of the 8 x 8 coefficient grid  
9 leading to significantly increased run-lengths. For indices with magnitudes greater than  
10 5, even a run-length as small as 3 will result in an escape sequence. As more indices  
11 (associated with largest magnitude coefficients) are preserved, the number of escape  
12 sequences continues to increase albeit a decreasing slope (i.e., a decreasing rate of  
13 generation) with each unit increase in the number of preserved AC coefficients per block.  
14 After a relatively small number of preserved indices (in the range from 5 to 10), the  
15 number of escape sequences starts declining with further increase in the number of  
16 preserved indices.

17 There are two mechanisms contributing to the observed relation of the escape  
18 sequence count with the number of preserved coefficients in each block. The first is the  
19 decrease in the magnitudes of the smallest indices which made it to the list of preserved  
20 indices due to increased nonzero coefficient allowance per block. This magnitude  
21 decrease, decreases the likelihood of their generating escape sequences. But even more  
22 important and influential than this observation is the fact that the inclusion of a steadily  
23 increasing number of nonzero indices leads to a denser population of the 8x8 coefficient

1 block, effectively breaking long runs of zeros between two large magnitude coefficients.  
2 This will lead to a shortening of the run-length components of the symbols to be encoded  
3 and hence a reduction in the number of escape sequences generated as well as an  
4 increased tendency towards using (run, level) symbols associated with shorter VLCs.

5 Even after the improvement achieved by the Pivot-2 procedure, there still remains  
6 quite a significant number of escape sequences. The effectiveness of and therefore the bit  
7 savings associated with the Pivot-2 procedure are limited since, more often than not,  
8 either there are no genuine candidate pivot indices between two qualifying largest  
9 magnitude coefficient indices or it is not feasible to use a genuine pivot index since it  
10 requires more encoding bits to include it owing to the locations and/or magnitudes of  
11 both the pivot and the qualifying largest magnitude coefficient. The nature of the  
12 coefficient selection implemented by the LMCS procedure and a very interesting human  
13 visual system masking behavior pertinent to DCT basis images open the way to another  
14 enhancement in the pivot index insertion framework.

15 The sensitivity of the human visual system to different DCT basis images is  
16 different. Here, the sensitivity is defined as the reciprocal of the smallest magnitude (i.e.,  
17 the threshold amplitude) of the basis image which enables its detection by human  
18 observers. This sensitivity is also dependent on various other factors. Among these are  
19 the viewing conditions such as the ambient luminance, and the display parameters such as  
20 the display luminance (the mean background intensity) and the visual resolution of the  
21 display (specified in terms of the display resolution and the viewing distance). However,  
22 more important for bit reduction purposes are the so called image-dependent factors

1 which model the mechanisms generated by the simultaneous presence of more than one  
2 basis image.

3 One very significant image-dependent factor influencing the detectability of DCT  
4 basis images is the effect of contrast masking, as described in Andrew B. Watson, Joshua  
5 A. Solomon, A. J. Ahumada Jr. and Alan Gale, "DCT Basis Function Visibility: Effects  
6 of Viewing Distance and Contrast Masking," in B. E. Rogowitz (Ed.), Human Vision,  
7 Visual Processing, and Digital Display IV (pp. 99-108). Bellingham, WA: SPIE, 1994.

8 This paper includes the following basic model of contrast masking:

$$m_T = t_T \max \left[ 1, \left| \frac{c_T}{t_T} \right|^{w_T} \right],$$

9 where:

10 the subscript  $T$  implies association with the spatial frequency  $T = (u,v)$ ,  $u,v = 0,1,\dots,7$ ,  
11 defined on the basis of the indices of the corresponding DCT coefficient;

12  $c_T$  is the given DCT coefficient;

13  $t_T$  is the corresponding absolute threshold;

14  $w_T$  is an exponent that lies between 0 and 1; and

15  $m_T$  is the masked threshold.

16

17 In the above equation,  $m_T$  defines the maximum extent of deviation from the  
18 coefficient's original value  $c_T$  which will not be detected by a typical human observer,  
19 when the correspondingly weighted basis image is displayed. It is easy to see from this  
20 model that, typically, sensitivity to quantization error in a particular DCT coefficient,  
21 decreases with the magnitude of that coefficient due to the increased masked threshold.

1 Note that this first order model of contrast masking describes the sensitivity to a  
2 particular coefficient's quantization error as being independent of the magnitudes of all  
3 the other coefficients except for the DC coefficient. However, there is evidence to the  
4 contrary, which indicates that sensitivity to a particular coefficient's quantization error is  
5 affected by the magnitudes of other coefficients (i.e., inter-coefficient contrast masking)  
6 as described through the following model:

$$m_T = t_T \max \left[ 1, \left| f[T, M] \frac{c_M}{t_T} \right|^w \right],$$

$$f[T, M] = \exp \left[ -\pi \|T - M\|^2 / \zeta_T^2 \right],$$

$$\zeta_T = \zeta \max[1, \|T\|].$$

7  
8 In the above model:

9 the subscript  $T$  implies association with the spatial frequency  $T = (u, v)$ ,  $u, v =$   
10  $0, 1, \dots, 7$ , defined based on the indices of the corresponding DCT coefficient;

11 the subscript  $M$  implies association with the spatial frequency  $M = (x, y)$ ,  $x, y =$

12  $0, 1, \dots, 7$ , defined based on the indices of the corresponding DCT coefficient;

13  $c_T$  is the DCT coefficient associated with the given test DCT basis image;

14  $t_T$  is  $c_T$ 's corresponding absolute threshold;

15  $c_M$  is the DCT coefficient associated with the given masking DCT basis image;

16  $w$  is an exponent that lies between 0 and 1;

17  $f[T, M]$  is a positive, frequency-dependent scaling factor; and

18  $m_T$  is the masked threshold for  $c_T$  due to the presence of  $c_M$ .

1

2 Observe that  $f[T, M]$  assumes its maximum value of 1 when  $T = M$ . In this case,  
3 this latter improved model reduces to the first order model described above.  $f[T, M]$   
4 reflects the sensitivity of  $c_T$ 's detection to the presence of a masking coefficient  $c_M$  at the  
5 frequency  $M$ . The larger  $f[T, M]$  is, the stronger is the masking influence. (The precise  
6 extent of masking generated by  $c_M$  is also dependent on the ratio  $c_M / t_T$ ). The second  
7 equation in the description of the improved model defines  $f[T, M]$  as a radially  
8 symmetric Gaussian function parameterized through a single parameter  $\zeta$  defined in the  
9 third equation of the same description. It is interesting to note that the bandwidth of  $f[T,$   
10  $M]$  increases in proportion to (the  $L^2$  norm of the spatial) frequency except for the DC  
11 coefficient. This, in particular implies a reduced sensitivity to (equivalently an easier  
12 masking of) high-frequency coefficients. The three fundamental parameters of the  
13 refined model, namely  $w$ ,  $t_T$  and  $\zeta$  are determined through a least squares estimation  
14 method on empirical data.

15 The relation of the contrast masking phenomenon to the first two generations of  
16 the pivot index methodology is easy to conclude. The qualified largest magnitude  
17 coefficients selected by the LMCS algorithm typically have the potential to generate a  
18 strong masking effect in their close vicinity in the frequency domain due to an increased  
19  $c_M / t_T$  ratio. Furthermore, the nature of both zigzag scans (i.e., achieving a slowly  
20 changing frequency content) lead to the insertion of pivot indices which are (most of the  
21 time) very close to their respective qualified LMCS coefficients in the frequency domain.  
22 This in return leads to a smaller value for the metric  $\|T - M\|$  in the exponent of  $f[T, M]$ ,  
23 increasing its value too. (In case of alternate zigzag scan, there are a few cases of

1 potentially having a somewhat larger frequency difference between the frequency of the  
2 preserved index and the frequency of the pivot index compared to the case of the default  
3 zigzag scan. Yet, these increased differences when they are realized, only weaken the  
4 extent of inter-coefficient contrast masking but do not eliminate masking altogether.)  
5 Given the two effects of large magnitude qualified coefficients generating a strong  
6 masking in their frequency domain neighborhoods and the pivot indices almost always  
7 being placed very close to the qualified coefficients, by moving from the first version  
8 (Pivot-1) to the second version (Pivot-2) of the pivot index methods, we achieve a further  
9 savings in the number of encoding bits at the expense of a marginal PSNR reduction  
10 which is visually masked.

11 We will now carry this idea one step further and consider introducing an artificial  
12 pivot when genuine ones are missing or inefficient for their purpose. Such an action will  
13 be taken only when it is beneficial to do so. For the best savings in the required number  
14 of encoding bits and for the least additional distortion, the pivot index should be a plus  
15 one or a minus one. There are several interesting observations and conclusions regarding  
16 this class of the pivot index technique and one of its possible implementations.

17 First, a pivot index should be placed in the position immediately preceding the  
18 position of the largest magnitude index in the adapted zigzag scanning order. This  
19 position is either the only position to get any savings or the position for the largest  
20 savings among all alternative positions. Consequently, a (run, level) symbol ( $n, M$ ) is  
21 transformed into the cascade combination of symbols  $(n-1, 1)$  and  $(0, M)$ . There are a  
22 few minor exceptions to this rule, and these exceptions can be encoded in a special table  
23 or tested for prior to a table lookup, as described below. A corollary to this observation is

1 that the implementation complexity is very low since there is no extensive decision and  
2 search processes involved as to where to place the pivot.

3 Second, the decision as to whether placing a pivot will help or not is as simple as  
4 a small table lookup. Given that we have to code the (run, level) symbol ( $n$ ,  $M$ ), we  
5 immediately know that the symbols ( $n-1$ , 1) and (0,  $M$ ) must be coded instead if a pivot is  
6 to be employed. The analysis of savings associated with this type of symbol  
7 transformation can be performed once off-line for all possible ( $n$ ,  $M$ ) pairs and a decision  
8 table to be indexed by  $n$  and  $M$ , can be generated to make the decision in the form of a  
9 Yes (1) or No (0) answer. For both of MPEG-2's VLC tables (Table 0 and Table 1) if the  
10 run-length ( $n$ ) is equal to 0 or greater than 32 or the level ( $M$ ) is greater than 40 in  
11 magnitude, the above proposed pivot technique cannot help. Therefore, the required  
12 table size is  $32 \times 40$  bits =160 Bytes.

13 Third, the degrading influence of such pivots on the reconstruction quality with  
14 respect to the PSNR metric is marginal since they are used in association with the LMCS  
15 approach. More importantly, as discussed above, the distortion introduced by the pivots  
16 is perceptually masked to a large extent due to inter-coefficient contrast masking.  
17 Moreover, as will be described below with reference to FIGs. 41 to 44, if a decoder is  
18 aware that the Pivot-3 technique is employed by an encoder (or a transcoder), then the  
19 decoder can, with a high accuracy, distinguish genuine indices with value (+/-) 1 from the  
20 inserted, noise-like pivot indices, and therefore the decoder can remove most of the  
21 inserted, noise-like pivot indices to avoid any significant degradation in PSNR.

22 With reference to FIG. 36, there is shown a sequence of DCT coefficients ... $C_i$   
23 ... $C_j C_k \dots$  in the coefficient scan order giving rise to a sequence of (run, level) symbols

1 for encoding an 8x8 block of pixels. In particular, the run length for the (run, level)  
2 symbol to be used for encoding the coefficient  $C_k$  is determined by the number R of  
3 consecutive AC coefficients having a zero level and immediately preceding the  
4 coefficient  $C_k$  in the scan order. In this example, the coefficients  $C_i$  and  $C_k$  are  
5 superscripted with asterisks to indicate that they have non-zero levels. The Pivot-3  
6 technique decides whether or not the (run, level) coding for each non-zero AC  
7 coefficient, such as the coefficient  $C_k$ , should be modified by changing the level of an  
8 immediately preceding coefficient  $C_j$ , in the scan order, from level 0 to a level of  
9 magnitude 1 (i.e., from 0 to a level of +1 or -1).

10 In order to decide whether or not the (run, level) coding for each non-zero AC  
11 coefficient should be modified by changing the level of an immediately preceding  
12 coefficient in the scan order, a lookup operation can be performed on a two-dimensional  
13 pivot table having a respective entry for each possible run length and level magnitude.  
14 As shown in FIG. 37, for example, the pivot table may include 64 rows for each possible  
15 encoded run length from 0 to 63, and 2048 columns for each possible encoded level  
16 magnitude from 1 to 2048. Inspection of such a pivot table, however, shows that no pivot  
17 should (or can) be inserted for a run length of zero, a run length greater than 32, or a level  
18 magnitude greater than 40. Therefore, considerable table memory may be saved by only  
19 storing a partial pivot table (497) having 32 rows and 40 columns.

20 With reference to FIG. 38, there is shown a first sheet of a flowchart for the Pivot-  
21 3 procedure. In a first step 501, scanning of a stream of indices begins in a next block. If  
22 the end of the stream is reached, as tested in step 502, then the procedure is finished.  
23 Otherwise, execution continues to step 503, to get the next index to be coded. If the end

1 of the block is reached, as tested in step 504, then execution loops back to step 501.  
2 Otherwise, execution continues to step 505 to determine the next (run, level) pair (M, N).  
3 In step 506, the run and level values are used to lookup a pivot table. If the pivot table  
4 indicates that a pivot should be inserted, then execution continues from step 507 to step  
5 508 in FIG. 39. Otherwise, execution branches from step 507 to step 513 in FIG. 39.

6 In step 508 of FIG. 39, a lookup is performed on the original index block for the  
7 immediately preceding location in the scan order to see if there was a genuine index to be  
8 discarded; in other words, an index for a non-zero, non-qualifying AC coefficient in the  
9 original-quality picture. If there is such an index, as tested in step 509, then execution  
10 continues to step 510 to lookup the VLC table for a run length of M-1 and a level equal to  
11 one in magnitude with a sign the same as the sign of the found index, in order to insert a  
12 pivot index having the corresponding VLC code. Then, in step 512, the VLC table is  
13 looked up for a run length of zero and a level of N, to re-code the variable-length code for  
14 the index obtained in step 503. In other words, instead of coding the variable-length code  
15 for the (run, level) of (M, N), a savings in bits is achieved by coding the variable-length  
16 code for (M-1, SIGN(INDEX)\*1) followed by the variable-length code for (0, N). From  
17 step 512, execution loops back to step 503 in FIG. 38. If there is no genuine non-zero,  
18 non-qualified index as tested in step 509, then execution continues to step 511 to lookup  
19 the VLC table for a run length of (M-1) and a level equal to one in magnitude with  
20 always a positive sign, in order to insert a pivot index having the corresponding VLC  
21 code. After step 511, the execution continues to step 512. In other words, instead of  
22 coding the variable-length code for the (run, level) of (M, N), a savings in bits is achieved  
23 by coding the variable-length code for (M-1, 1) followed by the variable-length code for

1 (0, N). It should be noted that for the artificial pivot index inserted in step 511 we  
2 arbitrarily but uniformly chose a positive sign, which to some extent supports the  
3 discrimination of artificial pivot indices from similar-looking genuine indices at the  
4 decoder. If a pivot is not to be inserted, then in step 513 of FIG. 39, the VLC table is  
5 looked-up for a run length of M and a level of N, in order to code the variable-length  
6 code for (M, N). Execution loops from step 513 back to step 503 in FIG. 38.

7 With reference to FIG. 40, there is shown a flowchart of a subroutine that  
8 simulates a lookup in the pivot table of FIG. 37 by performing some comparisons and, if  
9 necessary, performing a lookup of the partial pivot table (497 in FIG. 37). In a first step  
10 521 in FIG. 40, if the run length is zero, then execution branches to step 522, to return an  
11 indication that a pivot is not to be inserted. If the run length is not zero, then execution  
12 continues from step 521 to step 523. In step 523, if the run length is greater than 32,  
13 execution branches to step 522, to return an indication that a pivot is not to be inserted. If  
14 the run length is not greater than 32, then execution continues from 523 to step 525. In  
15 step 525, the level magnitude is computed as the absolute value of the level. In step 526,  
16 if the magnitude is greater than 40, then execution branches to step 522 to return an  
17 indication than a pivot is not to be inserted. In step 526, if the magnitude is not greater  
18 than 40, execution continues to step 527. In step 527, a lookup is performed upon the  
19 partial pivot table.

20 Following is a listing of the partial pivot table, for VLC coding according to the  
21 MPEG-2 VLC coding Table One:

22

23 % Partial pivot table description.

1 % value 0 => Pivot is not to be used. Code as it is.

2 % value -1 => Use a pivot only if there is already a non-zero coefficient in the

3 pivot position.

4 % value of 1 => Use a pivot.

5 % "zeros(n)" denotes a row vector of "n" zeros; "ones(n)" denotes a row vector of

6 "n" ones.

7

8 zeros(2) -1 -1 -1 ones(10) zeros(3) ones(22):

9 zeros(1) -1 -1 ones(37);

10 zeros(2) ones(38);

11 zeros(2) ones(38);

12 zeros(2) ones(38);

13 zeros(2) ones(38);

14 0 1 ones(29) zeros(9);

15 0 1 ones(29) zeros(9);

16 0 1 ones(29) zeros(9);

17 0 1 ones(29) zeros(9);

18 0 1 ones(29) zeros(9);

19 0 1 ones(13) zeros(25);

20 0 1 ones(13) zeros(25);

21 0 1 ones(13) zeros(25);

22 0 1 ones(13) zeros(25);

23 0 1 ones(13) zeros(25);

1           0 1 ones(13) zeros(25);  
2           0 1 ones(13) zeros(25);  
3           0 1 ones(13) zeros(25);  
4           0 1 ones(13) zeros(25);  
5           0 1 ones(13) zeros(25);  
6           0 1 ones(13) zeros(25);  
7           0 1 ones(13) zeros(25);  
8           0 1 ones(13) zeros(25);  
9           0 1 ones(13) zeros(25);  
10          0 1 ones(13) zeros(25);  
11          0 1 ones(13) zeros(25);  
12          0 1 ones(3) zeros(35);  
13          0 1 ones(3) zeros(35);  
14          0 1 ones(3) zeros(35);  
15          0 1 ones(3) zeros(35);  
16          0 1 ones(3) zeros(35);

17

18         The partial pivot table listed above has a few entries of value -1. For each of  
19         these entries, the insertion of a pivot results in exactly the same number of bits that would  
20         be needed if a pivot were not inserted. Therefore, in these cases a pivot should be  
21         inserted only if there is a genuine nonzero index (i.e., an index having a non-zero level in  
22         the original picture) at the pivot location. Pivot inclusion in the location of such an  
23         already existing nonzero index will help to improve the signal quality without using any

1 more bits for (run, level) encoding. However, when such a pivot insertion is made, in  
2 accordance with the general pivot insertion rules, even if the level magnitude in the  
3 original picture is greater than one, the level magnitude of the inserted pivot should be set  
4 to one, and the sign of the level of the inserted pivot should be the same as the sign of the  
5 level in the original picture.

6 Storage of the -1 values in the memory allocated to the table would unduly  
7 increase the amount of memory needed for the table. Instead, these few entries can be  
8 coded in the table look-up process as follows:

9

10           % lookup of the partial pivot table

11           START       IF (RUN = 1) THEN GOTO 100

12                       IF (RUN = 2) THEN GOTO 200

13           50           PPTV <= PPT(RUN, MAG)

14                       RETURN

15           % table lookup returns a value of PPTV = 0 or 1

16           100          IF (MAG<3) THEN GOTO 50

17                       IF (MAG>5) THEN GOTO 50

18           150          PPTV <= -1

19                       RETURN

20           % returns a value of PPTV = -1

21           200          IF (MAG<2) THEN GOTO 50

22                       IF (MAG>3) THEN GOTO 50

23                       GOTO 150

1  
2       Returning now to FIG. 40, in step 528, if the partial pivot table value (PPTV), for  
3       the row = RUN and column = MAG, has a value of zero, then execution branches to step  
4       522 to return an indication that no pivot should be inserted. Otherwise, execution  
5       continues from step 528 to step 529. In step 529, if the partial pivot table value (PPTV)  
6       is equal to 1, then execution branches to step 524 to return an indication that a pivot  
7       should be inserted. Otherwise, for the case of PPTV = -1, execution continues to step  
8       530. In step 530, execution branches depending on whether or not there is a non-  
9       qualifying coefficient (i.e., an AC coefficient having a non-zero level in the original  
10      picture) in the pivot position. If so, then execution branches to step 524 to return an  
11      indication that a pivot should be inserted. If not, then execution branches to step 522 to  
12      return an indication that a pivot should not be inserted.

13      FIG. 41 shows the Pivot-3 method of inserting pivot indices during the encoding  
14      or transcoding process and partial removal of the pivot indices during the decoding  
15      process. Transform coefficients are obtained from an original block-coded picture 541.  
16      The encoding or transcoding process includes the Pivot-3 method 542 of inserting noise,  
17      in the form of pivot indices, to reduce the number of bits for (run, level) coding of the  
18      transform coefficients. The reduction in the number of bits facilitates the transmission or  
19      storage 543 of the (run, level) coded transform coefficients. The decoding process 544  
20      includes the partial removal of the noise (i.e., the pivot indices) by removal of possible  
21      pivot indices not likely to occur in the original block-coded picture. The resulting  
22      transform coefficients are then used in a decoding process 545 of producing a  
23      reconstructed block-coded picture.

1        In order to facilitate the removal of possible pivot indices not likely to occur in an  
2 original block-coded picture, the Pivot-3 encoding method can be adjusted depending on  
3 whether or not the decoder will attempt removal of pivots. For example, as shown in  
4 FIG. 42, if the decoder will not attempt removal of pivots, as tested in step 551, then the  
5 process of encoding or transcoding will insert artificial pivots having a level of +1 or -1  
6 selected in a substantially random fashion, as shown in step 552. (For example, step 511  
7 of FIG. 39 would be modified to lookup the VLC table for either (M-1, 1) or (M-1, -1)  
8 selected by a pseudo-random number generator function.) However, if the decoder will  
9 attempt removal of pivots, then the level (more precisely the sign) of the artificial pivot  
10 indices should be selected in a convenient way such that the decoder will know whether  
11 or not a +1 or -1 is selected for the level of any pivot index inserted into the (run, level)  
12 coding. The most convenient way to perform such a selection of the pivot index level is  
13 to insert pivot indices all having the same level of either +1 or -1, such as a level of +1 as  
14 shown in step 553. (This is also what is shown in step 511 of FIG. 39.) Therefore,  
15 during the decoding process, all indices with value -1 (on average forming 50% of the  
16 genuine indices with magnitude 1) are certainly known to be genuine indices that should  
17 not be removed.

18        For distinguishing the genuine indices from the pivots within the set of indices  
19 with value +1, one can apply several rules. A (+1) level index immediately followed by a  
20 (+/-) 1 level index or an end-of-block (EOB) symbol in the scan order is a genuine index.  
21 A (+1) level index which is not immediately followed by another nonzero index in the  
22 scan order is a genuine index. If a (+1) level index and the nonzero index immediately  
23 following it together form an (n, +1) and (0, M) symbol pair, (i.e., a potential pivot

1 location encountered), still many cases exist in which depending on the values of the  
2 tuple ( $n$ ,  $M$ ), we can identify with certainty whether that (+1) level index is a genuine  
3 index or a pivot. For example, if  $M = 32$  and  $n > 6$ , we know with certainty that the (+1)  
4 level index is not a pivot index. That is,  $(7, 32)$ ,  $(8, 32)$ , ... $(31, 32)$  are the set of tuples  
5 for which the (+1) level index is a genuine index. It is because for these tuples the total  
6 number of bits to code both  $(n, +1)$  and  $(0, 32)$  is not less than 24 bits required to code  
7  $(n+1, M)$ .

8 FIG. 43 shows in greater detail the procedure of attempted pivot removal during  
9 decoding. In a first step 561, the decoder gets the next (run, level) pair. In step 562, if  
10 the end of the current encoded block is reached, then the procedure is finished.  
11 Otherwise, execution continues to step 563. In step 563, if the coefficient encoded by the  
12 (run, level) pair is not possibly a pivot, then execution branches to step 565 to accept the  
13 coefficient. Otherwise, if the coefficient is possibly a pivot inserted by the Pivot-3  
14 technique, then execution continues to step 564. In step 564, if the coefficient is not  
15 likely to be a pivot inserted by the Pivot-3 technique, then execution branches to step 565  
16 to accept the coefficient. Otherwise, if the coefficient is likely to be a pivot inserted by  
17 the Pivot-3 technique, then execution continues to step 566 to reject the coefficient. After  
18 step 565 or 566, execution loops back to step 561 to process the next (run, level) pair.

19 FIG. 44 shows further details of the process of determining whether or not an  
20 index is possibly a pivot (corresponding to step 563 in FIG. 43) and whether or not an  
21 index is likely to be a pivot (corresponding to step 564 in step FIG. 43). In a first step  
22 571 of FIG. 44, if the level is not equal to 1, execution branches to accept the coefficient.  
23 Otherwise, execution continues to step 572. In step 572, the decoding process looks

1 ahead to the immediately following symbol in the (run, level) symbol stream. If this  
2 immediately following symbol is an end-of-block (EOB) symbol as tested in step 573,  
3 then execution branches to accept the coefficient. Otherwise, execution continues to step  
4 574. In step 574, if the run length of the immediately following symbol is not zero, then  
5 execution branches to accept the coefficient. Otherwise, execution continues to step 575.  
6 In step 575, the magnitude of the level of the immediately following symbol is computed.  
7 Then in step 576, if the magnitude of the immediately following symbol is not greater  
8 than one, execution branches to accept the coefficient. Otherwise, execution continues  
9 from step 576 to step 577. In step 577, the decoder computes the run length that the  
10 immediately following symbol would have had if the coefficient (i.e., the possible pivot)  
11 were rejected. This is done by incrementing the run length by one. Then in step 578, the  
12 decoder looks up the pivot table with the run length (from step 577) and the magnitude  
13 (from step 575) in order to determine what the encoder would have done if the coefficient  
14 had a zero level in the original picture. In step 579, if the encoder would not have  
15 inserted the coefficient as a pivot, then execution branches to accept the coefficient.  
16 Otherwise, if the encoder would have inserted the coefficient as a pivot, then the  
17 coefficient is rejected as it is likely to have been inserted by the encoder.

18 It should be noted that steps 577, 578 and 579 could be omitted, in order to reject  
19 the coefficient if step 576 finds that the magnitude of the immediately following symbol  
20 has a level magnitude greater than one. Since the possible adverse effects of the pivots  
21 on the perceived picture quality is so small, it may not be worthwhile to perform steps  
22 577 to 579. However, if the table memory is allocated anyway for encoding purposes, for

1 example in a transceiver application or a picture storage and retrieval application, then  
2 the cost of performing steps 577, 578, and 579 would be minimal.

3 It should also be noted that the pivot table used in step 578 could be slightly  
4 different from the pivot table used for encoding, in order that the pivot table for decoding  
5 and rejecting possible pivots could take into account the probability that, when step 578  
6 is reached, a (RUN, MAG) pair would occur in the original picture, and the pivot table in  
7 step 578 would indicate the insertion of a pivot point, yet a pivot would not have been  
8 inserted by the encoder or transcoder because the original picture would include an  
9 immediately preceding coefficient of level = 1. Such slight differences in the tables  
10 could occur for (RUN, MAG) pairs having both a short run length and a small magnitude,  
11 and they could be found by encoding a series of test pictures and computing a histogram  
12 indicating, for each possible (RUN, MAG) pair, the percentage of the time that the  
13 procedure of FIG. 44 rejects a coefficient found in the original picture (run, level) coding  
14 when steps 577 to 579 are reached. If this percentage is greater than 50%, then the pivot  
15 table entry for the (RUN, MAG) pair should be changed so that this percentage would  
16 become less than 50%.

17 Although the pivot insertion procedures provide the most significant  
18 improvements when used in conjunction with the LMCS scaling procedure, the pivot  
19 insertion procedures may also provide substantial reductions in the bits required for  
20 encoding when used together with the LMIS scaling procedure or other scaling or  
21 encoding procedures. FIG. 45, for example, shows a flow diagram for using both the  
22 LCMS procedure and the LMIS procedure in combination with the pivot insertion  
23 techniques for both transcoding and encoding. For transcoding, a (run, level) encoded bit

1 stream for a picture is produced by a high resolution encoder 581 or a low resolution  
2 encoder 582. A decoder 583 decodes the (run, level) encoded bit stream.

3 For using the LMCS procedure, a de-quantizer 584 de-quantizes the levels to  
4 produce corresponding coefficient values. Coefficient selection 585 is performed  
5 according to the LMCS procedure, and pivot insertion 586 to reduce the number of bits  
6 for encoding is performed on the selected coefficients, to produce a (run, level) encoded  
7 picture.

8 Coefficient selection 587 by the LMIS procedure is performed on the decoded  
9 (run, level) information, and the pivot insertion 586 is performed on the selected  
10 coefficients to produce a (run, level) encoded picture.

11 When using the LMCS or LMIS procedures with pivot insertion during encoding,  
12 a DCT encoder 588 produces DCT coefficient values for 8x8 pixel blocks in the picture.  
13 Coefficient selection 589 by LMCS is followed by coefficient quantization to produce  
14 corresponding level values, which are used during the pivot insertion 586 to produce the  
15 (run, level) encoded picture. For the LMIS procedure, the DCT coefficient values from  
16 the DCT encoder 588 are processed by a quantizer 591 to produce a series of  
17 corresponding level values, and coefficient selection 592 by LMIS produces a subset of  
18 these level values for the pivot insertion 586 to produce the (run, level) encoded picture.  
19 The 8x8 bank of quantizers 591 may have quantization step sizes that are not uniform  
20 within the 8x8 block of DCT coefficients, for example, so that the higher frequency  
21 coefficients in each block are quantized with a larger step size than the lower frequency  
22 coefficients in the block.

1        Although the LMCS, LMIS, and pivot insertion methods have been described  
2        with respect to reducing the number of bits for encoding pictures which are MPEG-2  
3        video frames, it should be understood that the methods have general applicability to  
4        reducing the number of bits for (run, level) encoding regardless of the information  
5        represented by coefficients that are encoded. For example, the methods are directly  
6        applicable to scaling and encoding of individual pictures encoded by JPEG, which also  
7        run-length encodes DCT coefficients for 8x8 pixel blocks. The methods are also  
8        applicable to the use of other transform encoding techniques such as Fourier transform or  
9        wavelet transform techniques, and the encoding and compression of one-dimensional  
10      signals, such as audio signals.